IEEE Signal Processing MAGAZINE
Volume 38 | Number 3 | May 2021

INNOVATION STARTS WITH EDUCATION

The Evolution of Signal Processing

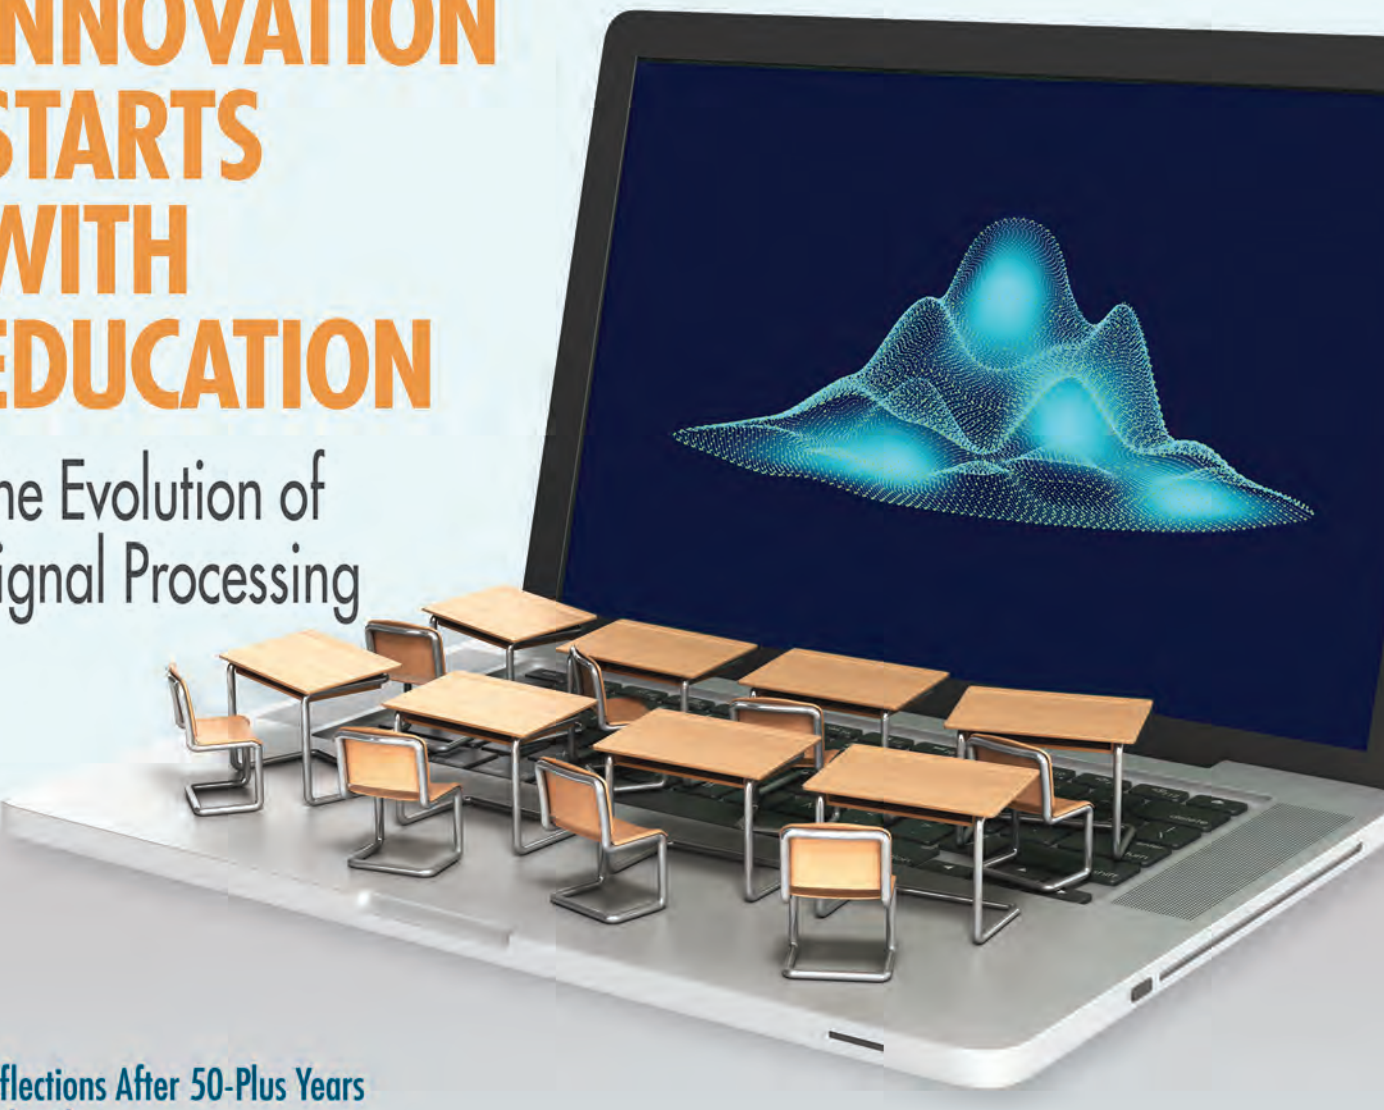

Reflections After 50-Plus Years in the Classroom

Proper Definition and Handling of Dirac Delta Functions

Alternative Data Paths for the Cascaded Integrator–Comb Decimator

IEEE Signal Processing Society

IEEE

CALL FOR PAPERS
55th Annual Asilomar Conference on Signals, Systems, and Computers

Asilomar Hotel and Conference Grounds
Pacific Grove, California
October 31, 2021 - November 3, 2021
www.asilomarssc.org

Authors are invited to submit papers in the following areas:

A. Communications Systems: 1. Modulation and Coding, 2. Cognitive Radios and Spectrum Sharing, 3. Ultra-Low Latency, 4. Optimization, 5. Physical Layer Security and Privacy, 6. mmWave, 7. Underwater, 8. Wireline and Optical Communications, 9. Satellite, 10. IoT, 11. V2V, 12. 5G and Beyond

B. MIMO Communications and Signal Processing: 1. Multiuser and Massive MIMO, 2. Channel Estimation and Equalization, 3. Full-Duplex, 4. Cooperation and Relaying, 5. Interference Management and Awareness, 6. mmWave, 7. Optimization

C. Networks: 1. Wireless Networks, 2. IoT, 3. Network Information Theory, 4. Optimization, 5. Graph Signal Processing, 6. Social Networks, 7. Distributed Algorithms, 8. Security, 9. Computational Offloading, 10. Self-Organizing Networks, 11. UAV and V2V Networks, 12. Smart Grid

D. Adaptive Systems, Machine Learning, Data Analytics: 1. Compressive Sensing, 2. Supervised Learning, 3. Unsupervised Learning, 4. Estimation and Inferences, 5. Adaptive and Cognitive Systems, 6. Adaptive Filtering, 7. Fast and Scalable Algorithms, 8. Models for High-Dimensional Large-Scale Data, 9. Distributed Computation and Storage, 10. Deep Learning, 11. Reinforcement Learning

E. Array Processing and Multi-Sensor Systems: 1. Source Localization and Separation, 2. Beamforming, 3. Robust Methods, 4. MIMO and Cognitive Radars, 5. Tensor Models and Processing, 6. Sparse Sensor Arrays, 7. Optimization, 8. Applications (Imaging, Radar, Sonar, Microphone Arrays)

F. Biomedical Signal and Image Processing: 1. Molecular and Medical Imaging, 2. Computational Imaging, 3. Neuroengineering, 4. Processing of Physiological Signals, 5. Bioinformatics and Computational Biology, 6. Image Registration and Multimodal Imaging, 7. Functional Imaging, 8. Brain Machine Interfaces, 9. Neural Signal Processing, 10. Visualization

G. Architectures and Implementation: 1. Energy Efficiency, 2. Accelerators, 3. Reconfigurable Processing, 4. Non-idealities, 5. Multicore, Many-core and Distributed Systems, 6. Algorithm and Architecture Co-optimization, 7. Architectures for Big Data, 8. Architectures for Machine Learning, 9. Cyber-Physical Systems, 10. Test-beds, 11. Mixed-Signal Processors, 12. Arithmetic and Algorithms

H. Speech, Image and Video Processing: 1. Speech Coding, 2. Speech Recognition, 3. Audio Coding, 4. Document Processing, 5. Models for Speech and Image Processing, 6. Image and Video Coding, 7. Learning and Autonomous Systems, 8. Computer Vision, Image and Video Analysis, 9. Image and Video Forensics, 10. Biometrics and Security, 11. Hybrid Imaging Systems

Submissions must include a 100–200 word abstract, an extended summary (500–1000 words plus figures), the title of the paper, each author's name and affiliation, technical area, and topic designation from the above list. To be eligible to take part in the student paper competition the submission of a full paper (5-8 pages) is required. The conference website (www.asilomarssc.org) provides specific information on the electronic submission and publication processes. The conference executive committee and IEEE reserve the right to exclude a final paper submission from inclusion in the conference proceedings, if the final paper submission does not meet conference or IEEE standards for scope and/or quality. Submissions will be accepted starting March 23, 2021. **All submissions must be received by May 1, 2021**.

- **No more than FOUR submissions are allowed per contributor, as author or co-author.**
- **To have a paper published, an author must register and present the paper at the conference.**

Notifications of acceptance will be mailed by mid-July 2021. Full papers are due shortly after the conference and published in early 2022. All technical questions should be directed to the Technical Program Chair **Prof. Mario Huemer** (mario.huemer@jku.at) or to the General Chair **Prof. Martin Haardt** (martin.haardt@tu-ilmenau.de). Prospective organizers of special sessions containing typically four papers are invited to submit proposals to the General or Technical Chair by January 31, 2021. Proposals must include title, topic, rationale, session outline, contact information, and a description of the session organization with paper titles and contact information for proposed special session authors.

CONFERENCE COMMITTEE (* serving in his/her capacity)

General Chair: Martin Haardt, *Ilmenau University of Technology*
Technical Program Chair: Mario Huemer, *Johannes Kepler University Linz*
Conference Coordinator: Monique P. Fargues*, *Naval Postgraduate School*
Publication Chair: Michael Matthews, *Northrop Grumman Innovation Systems*
Publicity Chair: Linda S. DeBrunner, *Florida State University*
Finance Chair: John D. Roth*, *Naval Postgraduate School*
Electronic Media Chair: Marios S. Pattichis, *University of New Mexico*

The site for the 2021 Conference is at the Asilomar Conference Grounds, in Pacific Grove, CA. The grounds border the Pacific Ocean and are close to Monterey, Carmel, and the scenic Seventeen Mile Drive in Pebble Beach. The Conference is organized by the non-profit Signals, Systems and Computers Conference Corporation.

Digital Object Identifier 10.1109/MSP.2021.3069699

Contents

Volume 38 | Number 3 | May 2021

SPECIAL SECTION

INNOVATION STARTS WITH EDUCATION

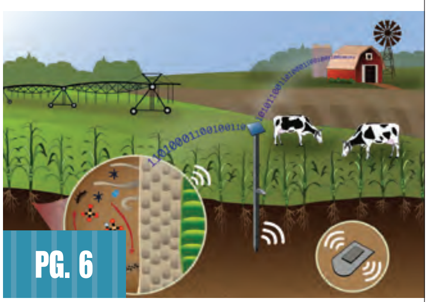

ON THE COVER

This special issue of *IEEE Signal Processing Magazine* includes 16 articles that revolve around three general and most pressing aspects of modern signal processing education: how to educate differently (better), student engagement, and the promotion of the societal impact of signal processing. Enjoy the issue!

COVER IMAGES—LAPTOP: ©SHUTTERSTOCK.COM/NOSOROGUA, GRAPHIC: ©SHUTTERSTOCK.COM/INGARA

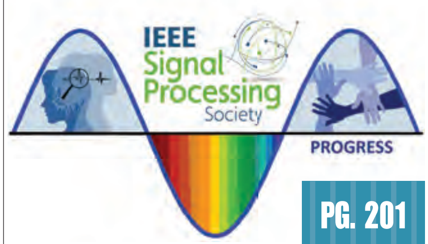

IEEE SIGNAL PROCESSING MAGAZINE (ISSN 1053-5888) (ISPREG) is published bimonthly by the Institute of Electrical and Electronics Engineers, Inc., 3 Park Avenue, 17th Floor, New York, NY 10016-5997 USA (+1 212 419 7900). Responsibility for the contents rests upon the authors and not the IEEE, the Society, or its members. Annual member subscriptions included in Society fee. Nonmember subscriptions available upon request. **Individual copies:** IEEE Members US$20.00 (first copy only), nonmembers US$248 per copy. Periodicals postage paid at New York, NY, and at additional mailing offices. **Postmaster:** Send address changes to IEEE Signal Processing Magazine, IEEE, 445 Hoes Lane, Piscataway, NJ 08854 USA. Canadian GST #125634188 **Printed in the U.S.A.**

Digital Object Identifier 10.1109/MSP.2021.3054855

COLUMNS

DEPARTMENTS

©SHUTTERSTOCK/DMITRI KALVAN

The IEEE Statistical Signal Processing Workshop will be held 11–14 July in Rio de Janeiro, Brazil.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

IEEE Signal Processing Magazine

Digital Object Identifier 10.1109/MSP.2021.3054856

SCOPE: *IEEE Signal Processing Magazine* publishes tutorial-style articles on signal processing research and applications as well as columns and forums on issues of interest. Its coverage ranges from fundamental principles to practical implementation, reflecting the multidimensional facets of interests and concerns of the community. Its mission is to bring up-to-date, emerging, and active technical developments, issues, and events to the research, educational, and professional communities. It is also the main Society communication platform addressing important issues concerning all members.

Christian Jutten | Editor-in-Chief | Christian.Jutten@grenoble-inp.fr

Introducing *SPM*'s New Team of Area Editors: Part 2

In my March editorial in *IEEE Signal Processing Magazine* (*SPM*) [1], I presented the new area editors for feature articles, outreach and social media, as well as for the eNewsletter. In this issue, I announce the two area editors for columns and forums and the new area editor for special issues. In addition to the area editors, as I previously explained in the March issue, *SPM* benefits from the valuable help and wide expertise, in signal and image processing (SIP) and its applications, of the senior and associate editors who form the Editorial Board.

Introducing three new area editors

Columns and forums

Rodrigo Guido

I have been in love with signal processing and electronics since my childhood, long before receiving my B.Sc. degrees in computer science and computer engineering. Since receiving my Ph.D. degree in signal processing from the University of São Paulo, Brazil, in 2003, I have been involved with editorial activities. As an associate professor at São Paulo State University, I have contributed to dozens of scientific journals, serving as a guest editor, an associate editor, and an area editor. I have focused my research on speech processing-related topics where, from the digital side, a combination of machine learning and wavelet transforms approaches are the main tools and, from the analog side, radio-frequency circuitry is a passion. I am extremely glad to serve as the co-area editor of columns and forums for *SPM*, contributing to our community of students, scientists, and engineers in the fantastic world of signal processing. Please do not hesitate to contact me at guido@ieee.org if you have any questions. It is always a great pleasure to help.

SPM's columns and forums articles are intended to provide high-level material focused on different content, types of coverage, and target audiences.

Vicky Zhao

I am an associate professor in the Department of Automation, Tsinghua University, China. Before joining Tsinghua University, I was with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, Canada, as an assistant professor (2006–2012) and associate professor (2012–2016). My research interests include media-sharing social networks, information security and forensics, digital signal processing, and communications. I received the IEEE Signal Processing Society 2008 Young Author Best Paper Award and the Asia-Pacific Signal and Information Processing Annual Summit and Conference 2020 Best Paper Award. I have been active in professional societies, serving as a senior area editor of *IEEE Signal Processing Letters* and an associate editor of *IEEE Open Journal of Signal Processing*, *IEEE Transactions on Information Forensics and Security*, *IEEE Signal Processing Letters*, and *SPM*. I also serve as an organizing committee member for IEEE conferences and workshops, including IEEE ICASSP, IEEE ICIP, IEEE International Conference on Multimedia and Expo, IEEE International Workshop on Information Forensics and Security, and IEEE International Workvshop on Multimedia Signal Processing. If you have any suggestions or would like to contribute, please feel free to contact me at vzhao@tsinghua.edu.cn.

Submit contributions

SPM's columns and forums articles are intended to provide high-level material focused on different content, types of coverage, and target audiences. In particular, they aim to deliver tutorial-like lessons with relevant insights and to track recent technological advances, balancing theoretical and experimental aspects and offering diversified coverage in the

Digital Object Identifier 10.1109/MSP.2021.3054857
Date of current version: 28 April 2021

wide signal processing area. As the area editors of *SPM*'s columns and forums, we encourage all researchers, engineers, and scientists to share relevant material, recent news, working experiences, views related to signal processing, and even funny stories and cartoons to lighten life a little bit. We look forward to receiving your contributions.

Special issues

Xiaoxiang Zhu

I received my M.Sc. degree, Dr.-Ing. degree, and "Habilitation" in the field of signal processing from the Technical University of Munich (TUM), in 2008, 2011, and 2013, respectively. I am a professor of data science in Earth observation at TUM, and I head the Department of Earth Observation Data Science at the German Aerospace Center. I am also the director of the international artificial intelligence (AI) future lab AI4EO; codirector of the Munich Data Science Institute, TUM; co-spokesperson for the Munich Data Science Research School; and head of the Helmholtz aeronautics, space, and transport AI research field. I was a guest scientist and visiting professor at the Italian National Research Council, Fudan University, University of Tokyo, and University of California, Los Angeles, in 2009, 2014, 2015, and 2016, respectively. My research interests include remote sensing and Earth observation, signal processing, machine learning, and data science, with a special application focus on monitoring global urbanization. I am a Fellow of IEEE and a member of the Junge Akademie/Junges Kolleg, Berlin–Brandenburg Academy of Sciences and Humanities, German National Academy of Sciences Leopoldina, and Bavarian Academy of Sciences and Humanities. *SPM* is one of my favorite magazines, and I am happy to serve as the area editor of special issues, which aim to address pressing signal processing research topics that are highly relevant to a broad community. Please do not hesitate to contact me at xiaoxiang.zhu@dlr.de if you wish to see a special issue on an interesting topic that has not been covered or organize a special issue yourself.

In this issue

The theme of this special issue is "Innovation Starts With Education." It is based on an ICASSP 2019 panel about education [2]. Clearly, this topic is of great interest since many of us are involved in teaching signal processing at the bachelor's, master's, and doctoral levels and as supervisors of trainees and postdocs at companies and in research labs. Sixteen articles constitute the special issue, proposing many approaches and telling us about various experiences at numerous universities worldwide.

But first, Prof. A. Oppenheim and Prof. A. Constantinides present their article "Reflections After 50-Plus Years in the Classroom" [3] in the "Reflections" column. I encourage you to read Prof. Oppenheim's previous columns about education [4] (in 1992) and research [5] (in 2006) to understand his views and how they have changed, especially in the context of technological advances but also during the COVID-19 pandemic. Basically, I take from these articles—I hope not to betray his thoughts—that in addition to instilling the basics in students, it is fundamental to be attentive to pupils' motivations and interests, to let them have some fun and develop their creativity. Another point concerns the dual role of mentor and friend: this is possible with Ph.D. students and postdocs because of the intensive interactions that are involved, but I believe it is much more difficult in bachelor's and master's classes that have tens of students. Finally, I note that Prof. Oppenheim never learned to teach, and surprisingly, it seems that this is true in most countries.

Beyond fostering SIP skills, enhancing creativity, and addressing multidisciplinary applications, projects open the way to honest benchmarking, which is the cornerstone of open-access and reproducible science.

In fact, all the articles in this special issue focus on some of these facets, which seem to be "invariants," and readers will profit from the authors' ideas, experiences, feedback, and reflections, which will certainly be helpful for developing and delivering lectures. I also discovered (after 40 years of teaching!) the concepts of Bloom's taxonomy and Kolb's cycle. One of the difficulties of teaching SIP is that it requires instructors to impart a mastery of mathematics and statistics. On the other hand, SIP is a dream discipline since it can be useful in so many interesting—and sometimes funny—applications. To develop students' creativity, and even to help them understand basic concepts, it is easy to choose music, speech, biomedical signals, images, robotics, and others that can be fun and stimulating. It is also clear that the authors select diverse applications that are strongly related to their own research activities, illustrating how teaching and research duality is essential. Of course, project-based learning is present in many approaches for enhancing students' creativity, and addressing problems to which there is no unique solution but many approaches that can be rigorously discussed, implemented, and evaluated.

Beyond fostering SIP skills, enhancing creativity, and addressing multidisciplinary applications, projects open the way to honest benchmarking, which is the cornerstone of open-access and reproducible science. Currently, with massive open online courses and virtual conferences, such as ICASSP and ICIP in 2020, there is a wide diversity of documents that can be used anywhere at any time and enhance more classical ways of teaching. Finally, there is intense pressure from students and industry to teach fashionable technologies such as machine learning and deep learning: for SIP instructors, this presents an opportunity to trade between white- and black-box methods—model and data driven—and make students attentive to the explainability and robustness of approaches that could seem like magic and that must be applied with critical thinking and intelligence.

Most of the experiments presented in this special issue have been conducted, at least partly, during the past year, which was a strange one for education. Practically overnight, we had to teach remotely, speaking to a screen, with very little interaction with our students. In a blink of an eye, practical work and group projects became more complicated, if not impossible. You can certainly relate through your own teaching experience: how you fared and what tricks you used to motivate your students and ensure their well-being as much as possible.

Practically overnight, we had to teach remotely, speaking to a screen, with very little interaction with our students.

I would like to mention that, beyond this special issue, in any *SPM* issue, you can share an interesting teaching experience in a "Lecture Notes" column, if possible, with additional materials, such as exercises and quizzes for student assessments and codes for experiments. If you are interested, feel free to email your ideas to the area editors for columns and forums, Prof. Guido and Prof. Zhao.

References

[1] C. Jutten, "Introducing SPM's new team of area editors: Part 1," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 3–5, Mar. 2021. doi: 10.1109/MSP.2020.3044239.

[2] V. Solo, M. S. Greco, D. Mandic, P. Djuric, A. Spanias, and M. F. Bugallo, "Innovation starts with education: ICASSP 2019 education panel [SP Forum]," *IEEE Signal Process. Mag.*, vol. 36, no. 5, pp. 135–147, Sept. 2019. doi: 10.1109/MSP.2019.2923280.

[3] A. Oppenheim and T. Constantinides, "Reflections after 50-plus years in the classroom," *IEEE Signal Process. Mag.*, vol. 38, no. 3, pp. 15–19, May 2021.

[4] A. V. Oppenheim, "A personal view of education," *IEEE Signal Process. Mag.*, vol. 9, no. 2, pp. 69–72, 1992. doi: 10.1109/79.127285.

[5] A. V. Oppenheim, "One plus one could equal three (and other favorite clichés)," *IEEE Signal Process. Mag.*, vol. 23, no. 6, pp. 10–12, 2006. doi: 10.1109/SP-M.2006.248707.

SP

John Edwards

Signal Processing Plays a Key Role in Environmental Research Projects

Keeping people and ecosystems alive and healthy is perhaps the 21st century's biggest challenge

Despite the impressive technological strides made over the years, human lives still depend very much on the natural environment. Fortunately, technology can now be used to help address critical environmental concerns in air quality, soil condition, and weather events. In all of these areas and many others, signal processing is supporting the ability to provide immediate and long-term observations and insights.

Efficient air-quality monitoring

It's generally accepted that the efficient monitoring of airborne particulate matter (PM), particularly particles with an aerodynamic diameter measuring less than 2.5 μm (PM2.5), is an important step toward sustaining and improving public health.

Acknowledging this fact, researchers at the Max Planck Institute for the Science of Light have developed a novel way to continuously monitor a local environment for both the size and optical properties of individual airborne particles. The technique utilizes optical forces to automatically capture airborne particles and then propel them into a hollow-core fiber where they can be studied and counted, providing a potentially better way to monitor air pollution levels.

On-the-fly particle metrology uses both advanced optics and signal processing to continuously monitor the size and refractive index of individual airborne particles in an open atmosphere, says research team leader Shangran Xie (Figure 1). "It can overcome several limitations of … existing methods, offering the ability of simultaneous measurement of particle size and refractive index, which can assist in identifying particle material, real-time particle metrology, highly reproducible results, and unlimited device lifetime," he explains.

Current commercially available particle counters are limited to counting

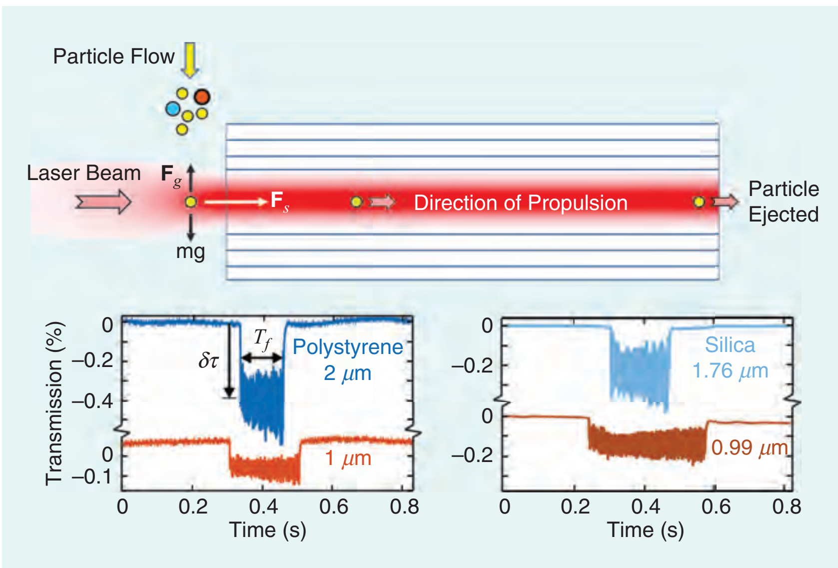

FIGURE 1. An on-the-fly particle metrology, developed by researchers at the Max Planck Institute for the Science of Light, uses both advanced optics and signal processing to continuously monitor the size and refractive index of individual airborne particles in an open atmosphere. (Source: Max Planck Institute for the Science of Light; used with permission.)

Digital Object Identifier 10.1109/MSP.2021.3058912
Date of current version: 28 April 2021

the number of airborne particles. If more detailed particle data are needed, the existing standard technique requires manual sampling with sophisticated equipment. While the combined approach can provide a full span of particle information, it's not a continuous measurement and can't provide real-time feedback on the pollutant.

The new technique promises to provide a reliable way to rapidly and continuously characterize airborne particles. "It can not only count the number of particles, which is related with the level of pollution, but also can provide detailed information on particle size distribution and chemical dispersion in real time," Xie says. "The configuration is also very simple; it's highly possible to build a shoe-box-sized device able to continuously monitor airborne PM2.5 particles in urban areas and industrial sites."

Xie reports that he and his team have been working on particle trapping and analysis in hollow-core photonic crystal fiber for years, gaining, over time, a deep understanding of particle scattering within a hollow-core fiber. "Inspired by the need for particle detection in air pollution monitoring, we think our fiber and the corresponding data processing procedure may offer a better solution," he says.

The new analysis approach traps airborne particles inside a laser beam by optical forces and propels them forward by radiation pressure. The trapping force is strong enough to overcome gravitational force on very small particles, such as PM2.5. The approach also automatically aligns the particles within the hollow-core fiber. Postalignment, laser light propels the particle into the fiber, forcing the laser light inside the fiber to scatter and create a detectable reduction in the fiber transmission.

At the project's heart is a novel signal processing algorithm that the researchers designed to retrieve useful information from the particle-scattering data in real time. A photodetector is used to convert the original optical signal into an electrical signal. "The goal of signal processing in our technique is to retrieve, as precisely as possible, the particle size and its refractive index from the fiber transmission signal," Xie states.

The fiber carries two types of information: transmission drop data created by particle scattering and time-of-flight information and the time it takes a particle to travel through the fiber. "Having these two [types of] information in hand, an algorithm based on particle scattering theory can be integrated into the signal processer to unambiguously retrieve the particle size and refractive index," Xie says.

A significant remaining signal processing challenge is dealing with a relatively weak transmission drop signal. "Normally, a single particle would only introduce a less than 1% signal drop; the signal we are [now] facing is a tiny drop lasting for tens of milliseconds on top of a strong dc background." Another concern is the particles that pass through the laser beam without being captured. "Those particles will introduce spikes in the signal which cannot be properly retrieved by the algorithm," Xie comments.

The biggest challenge the team now faces is finding an algorithm that can further translate particle number, size, and refractive index into PM2.5 concentration data as well as a description of the types of pollutants detected. "To do this, advanced signal processing algorithms on data classification … may be required to quickly identify the pollution type based on the known database," Xie says. "In other words, we need to further bridge the gap between the data in the lab and information for the end users."

Looking forward, Xie is hoping to further advance the system's particle characterization. "For instance, it may be possible to monitor the particle shape or surface roughness by analyzing the scattering patterns from the fiber endface or from the side," he says. "This can give additional information on the residence time of pollutants in the environment."

Improving tornado detection and tracking

University of Mississippi researchers believe that "listening" to tornadoes via infrasound will lead to significantly earlier and more accurate tornado warnings.

Despite the rapid advances in meteorological technology, detecting and tracking tornadoes remains a formidable task. More than 24,400 tornadoes have been reported across the United States since 2000, according to the National Centers for Environmental Information. Over the same period, tornadoes have killed almost 1,500 people and resulted in billions of dollars in damage.

Radar is unable to reliably detect tornadoes, states Roger Waxler, a University of Mississippi research associate professor of physics and astronomy and a principal scientist at the National Center for Physical Acoustics (NCPA). The wavelengths are too long and upward looking to generate accurate reports, he explains. Therefore, tornado warnings are currently issued solely on the basis of visual observations and/or confirmations.

Addressing this issue, significant effort has been poured into the development of short wavelength radar systems that might be able to detect tornadoes directly, perhaps by the debris generated by the tornado's funnel. "But these would require line-of-sight and would be blocked by hills, tree cover, and so on," Waxler observes.

Acoustics promise a better approach. Since sound doesn't depend on line-of-sight, it can detect tornado activity directly. An added benefit is that acoustic technology is generally less costly than radar systems, Waxler notes. "Tracking from acoustic technologies could assist in providing better estimates of locations to investigate and tornado passage times," he adds.

Joining Waxler in investigating acoustics' potential to detect and track tornados is Garth Frazier, a senior research scientist at NCPA and a University of Mississippi research associate professor of electrical engineering. Another key team member is Carrick Talmadge, also a senior NCPA research scientist and a University of Mississippi research associate professor of communication sciences and disorders.

For the past several years, the team has explored the potential of infrasound arrays that incorporate anywhere from five to 10 sensors. The sensors are installed directly on the ground in semi-permanent locations, measuring approximately 50 x 50 meters, that change on a seasonal basis (Figure 2). "We deploy a network of arrays based on guidance from meteorologists with the goal of covering a regional area," Frazier says.

In the current research phase, data are continuously logged at 1,000 samples/s by the Coordinated Universal Time-synchronized sensors, which run on solar-augmented battery power. "Periodically, the data are downloaded from the sensors during site visits, but this might be only once per several months," Frazier notes.

Using storm report information available on U.S. National Oceanic and Atmospheric Administration websites and the Iowa Environmental Mesonet website at Iowa State University, selected time periods of data are analyzed using array processing algorithms. "Signal processing is one of the three pillars of the technology, with the other two being infrasound sensors and long-range atmospheric sound propagation modeling and prediction," Frazier says.

The collected data are typically decimated to a lower rate, such as 100 Hz, prior to array processing. "Additionally, we high-pass filter to remove most signal fluctuations below 0.1 Hz," Frazier notes. The project uses three algorithms to estimate the directions of arrival and measured signal levels: a maximum likelihood approach based on the complex Wishart distribution, a signal subspace approach, and multiple signal classification. "All of the array signal processing is performed in the frequency domain," he adds.

The researchers selected the three specific signal processing approaches for their ability to resolve multiple sources simultaneously along different azimuths. "In two cases in Alabama, we have followed two tornadoes simultaneously," Frazier says. "In addition to storm-generated infrasound, we have to contend with anthropomorphic infrasound, especially from urban areas and industrial plants in particular."

During the course of their research, the team found that tornadic storms produce sound in various frequency bands. "Almost all of the long-range detections—50–100 km—we have observed have been in the band from 1 Hz to 10 Hz, and most of those have been in 2 Hz to 5 Hz," Frazier reports.

Much work remains to be done before that system can be used to generate reliable tornado warnings. "From a signal processing point of view, we still need to automate the entire processing pipeline into a real-time framework," Frazier says. Currently, all data processing is performed offline in either Python or Octave.

FIGURE 2. Tornado-detecting infrasound arrays, developed by University of Mississippi researchers, incorporate anywhere from five to 10 sensors. (Source: Shea Stewart/University of Mississippi; used with permission.)

The primary challenge the team now faces is frequent low signal-to-noise ratio issues. Wind noise created by intrinsic turbulent pressure fluctuations in the atmosphere surface layer combined with the interaction of wind and the sensor housing has been a particularly nagging concern.

The researchers also still need to perfect their real-time, bearings-only, 2D multiple-target tracking algorithms to provide accurate geolocations when multiple arrays detect and follow the same tornadoes. This task promises to be particularly challenging given the significantly different time delay that exists between sound emission and measurement at different sensors as the tornado travels along its path.

Another complicating factor is long-range atmospheric propagation, which can cause the average speed of sound between the source and receiver to vary when detected from different directions. "We have solved this problem for the single target case using a Bayesian framework that updates the state vector at the previously estimated emission time then propagates forward to the newly estimated emission time—the reverse of Kalman filtering steps," Frazier explains. "These calculations require the use of a ray-tracing model for the sound propagation that depends on a local vertical profile model of wind and temperature."

Waxler believes that the technology has reached the stage where it's time to begin moving toward small-scale demonstration test systems, including data telemetry for remote processing. "There's a clear path to the signal processing implementations that are still required," he says. "There will possibly be hiccups in the data transmission process, but the only way to address these is to begin to gain experience."

Sensor promises larger crop yields using less fertilizer

Soil ecosystems provide most of the antibiotics used to combat diseases, control the movement of water and chemical substances between the Earth and its atmosphere, and function as the source and storage media for gases that are critical to sustaining life, such as oxygen, carbon dioxide, and methane.

Facing rapid global population growth, climate change, and increasing competition for land resources, there's an urgent need to find a way to quickly and efficiently analyze and monitor soils. Being able to rapidly and reliably measure levels of soil phosphate—a finite and nonrenewable resource that stays in a complex biogeochemical environment—is particularly urgent since there have been calls for a global effort to utilize phosphate fertilizers as efficiently as possible.

Researchers at Kansas State University's industrial engineering department, working with their counterparts at educational institutions worldwide, are developing graphene sensor-based systems that can map and monitor soil phosphate levels while generating insightful real-time data. The project aims to help researchers and farmers understand soils better while increasing crop yields and minimizing the use of phosphate fertilizers (Figure 3).

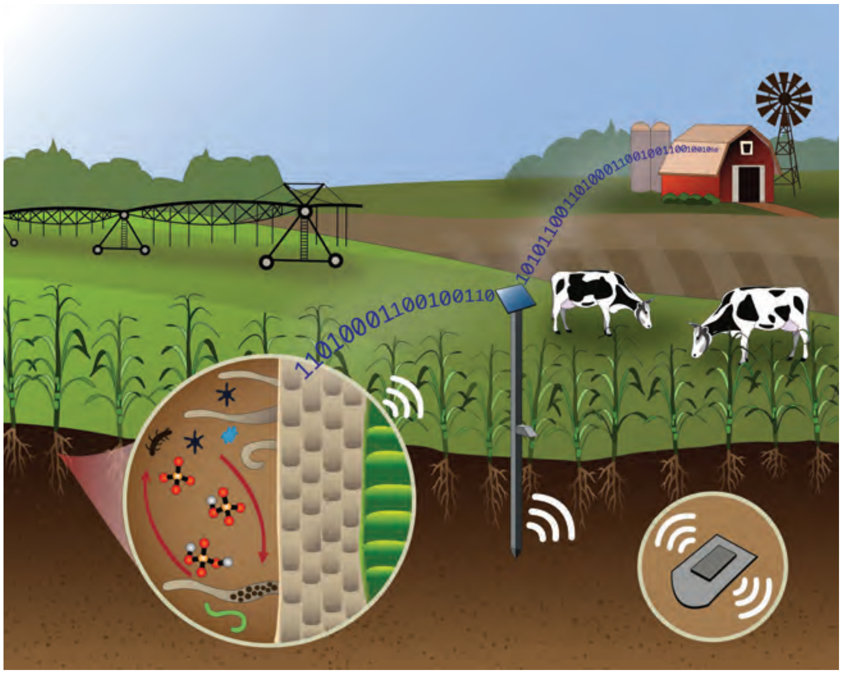

FIGURE 3. Printed graphene electrochemical sensors, combined with pH sensors and soil moisture sensor arrays on mechanically flexible substrates, deployed in the soil for phosphate sensing. A wireless circuit in combination with a custom-built impedance analyzer will transmit the data to the acquisition center. (Source: Suprem Das/Kansas State University; used with permission.)

While phosphate is a key crop nutrient, it's currently difficult for farmers to quickly and reliably map soil phosphate content levels since the process requires sending samples to a lab. Mapping and monitoring soil with portable and affordable sensors promise to provide a more accurate understanding of how soil composition changes over time, helping farmers to apply phosphate fertilizers only to the areas where it's most needed.

The project's researchers are focusing their efforts on two major areas: a soil sensor made of graphene (an atomically thin 2D carbon material) and a hardware-supported signal processing architecture. "Exploiting the interaction of phosphates with graphene will produce a characteristic signal read by an impedance that will be carefully collected and processed by the signal processing hardware," explains principal investigator Suprem Das, an assistant professor in industrial manufacturing systems engineering at Kansas State University's Carl R. Ice College of Engineering. "Therefore, identifying the fundamental impedance related to phosphates as well as discriminating signals from interfering species in the soil during the signal processing are important parts of our research," he says.

Signal processing plays a dual role in the project, which is funded by the U.S. National Science Foundation and U.K. Research and Innovation. "First, electrical signals from the soil collected by the sensor board need to be carefully analyzed to get the accurate estimate of phosphate content, so in situ processing of the signals is very important," says project co-investigator Biswajit Ray, an assistant professor of electrical and computer engineering at the University of Alabama in Huntsville. "Second, the phosphate content value needs to be transmitted wirelessly with limited power sources and [on a] resource-constrained hardware platform, so careful hardware design will be important," he adds.

The researchers plan to present data in a way that will allow end users to focus on the task at hand while ignoring the complex science underpinning the technology. "This technology is primarily aimed at the farming industry so that we can achieve a more sustainable agriculture," explains Adrien Chauvet, a lecturer in physical chemistry at the University of Sheffield as well as the project's primary United Kingdom investigator. "Such a technology would allow farmers to literally map the phosphate content of their crops, with an approximate square-mile resolution, live." The researchers envision a field deployment that includes multiple sensor boards distributed over a large area. Each board will be capable of measuring local phosphate content and transmitting the information to a control station.

Chauvet is confident that the technology has a promising future. "As a scientist, I see this project as the first step of a long-lasting collaboration that will go beyond the creation of the actual device," he states. "If we can prove that this sensing strategy works, then we can expand it and apply it to other minerals and heavy atoms."

Author

John Edwards (jedwards@johnedwardsmedia.com) is a technology writer based in Gilbert, Arizona, 85234, USA. Follow him on Twitter @TechJohnEdwards.

SP

Hagit Messer, Arye Nehorai, Jeff Krolik, José M.F. Moura,
Al Hero, and Joseph Tabrikian

In Remembrance of Peter Schultheiss

On 23 February 2021, Peter Schultheiss died peacefully in his sleep in Hamden, Connecticut, at the age of 96. He had been a faculty member at Yale University since 1947, where he taught for more than 65 years and trained dozens of Ph.D. students, postdocs, and research collaborators. His research interests centered on problems in detection and estimation theory. He led the most significant contributions to the theory of source localization by an array of sensors and advanced its application in underwater acoustics.

Prof. Schultheiss's research was motivated by real-world problems with direct connections to the physical world. He identified many of these through collaborations with U.S. Navy research laboratories. He insisted on accurate modeling and accepted approximations only when they were well justified, and he introduced a systematic method for discovering insights through the use of theoretical performance bounds. Invariably, modeling imperfections, such as limited coherence and calibration, constrain the ability of sensing systems to achieve optimal performance. By including uncertainty in his models, Prof. Schultheiss provided benchmarks to which the performance of real systems could be realistically compared. Moreover, his insights focused attention on areas of system design that most critically limited performance. His early application of the Cramér–Rao lower bound, from statistics, to sensor array systems spawned its use as well as that of other bounds in a wide variety of signal processing applications.

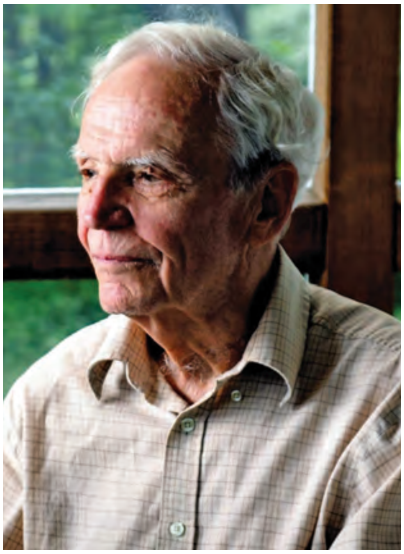

Prof. Peter Schultheiss.

More specifically, Prof. Schultheiss was an early and consistent contributor to passive and active source localization and time delay estimation in underwater acoustics and sonar environments. Working with his many students and collaborators, he derived optimal and suboptimal estimators and developed Cramér–Rao bounds that established limits on expected performance under a variety of conditions, including randomly perturbed arrays, the presence of interference, multipath, and unknown noise statistics. One such example is his pioneering work on array shape calibration using sources in unknown locations, under which he first presented the concept of a hybrid Cramér–Rao bound and demonstrated how its analysis could be used for studying the inherent limitations in practical parameter estimation problems.

> **He led the most significant contributions to the theory of source localization by an array of sensors and advanced its application in underwater acoustics.**

Prof. Schultheiss was highly recognized as an excellent teacher. He taught linear algebra in the Department of Mathematics for decades, as students preferred his teaching style and clarity. For his graduate students, postdocs, and research collaborators, he was an excellent mentor and was involved in the details of their research projects. He was a great inspiration to many researchers and his former students. Prof. Schultheiss was a kind individual who was always open to support younger researchers and provide them with valuable career guidance. He was highly appreciated by colleagues in the signal processing community and will be missed by all of them.

Digital Object Identifier 10.1109/MSP.2021.3065892
Date of current version: 28 April 2021

Mónica F. Bugallo, Anthony G. Constantinides, Danilo P. Mandic,
Alan V. Oppenheim, and Roberto B. Togneri

Innovation Starts With Education

Signal processing (SP) is at the very heart of our digital lives, owing to its role as the pivotal enabling technology for advancement across multiple disciplines. Its prominence in modern data science has created a necessity to supply industry, government labs, and academia with graduates who possess relevant SP expertise and are well equipped to deal with the manifold challenges in current and future applications. To this end, the ways to deliver both educational content and the core SP curriculum need to be revisited and integrated into current electrical engineering and computer science degrees to provide high-quality and hands-on multidisciplinary skills, experience, and inspiration for students at all levels.

SP education in today's universities is largely influenced by three modern trends:

1) the availability of competing and complementary online and multimedia resources
2) the fact that we live in a world in which the amount and diversity of information we generate, process, and analyze are growing
3) the explosive growth of computing power and the rapid development of new technologies for implementing both analog and digital SP.

These trends offer both opportunities and challenges, which we can and must exploit in charting dynamically adjustable courses that attract a high level of student engagement while offering a mix of essential background physics, intuition, mathematical rigor, and practical applicability of the taught material.

With such initiatives underway worldwide, this special issue aims to facilitate both keeping abreast with SP education and exploring innovative and participatory ways to present the educational materials. In effect, we cannot assume that students will be able to appreciate the scope and relevance of their courses without explicitly building a bridge between the material presented in class and cutting-edge research and the societal and practical impact of their education. This includes the convergence of educational material with other disciplines (machine learning, data science, big data, bioengineering, artificial intelligence, finance, and many others).

> "The role of a magician is to make simple things appear mysterious. The role of a teacher is to make mysterious things appear simple."
> —Al Oppenheim

This special issue of *IEEE Signal Processing Magazine* (*SPM*) therefore revolves around three general and most-pressing aspects of modern SP education:

- *How to educate differently (better)*: This includes the use of available technology, bringing research into the classroom, web resources, experiential learning, and massive open online courses (MOOCs).
- *Student engagement*: This includes ways to enhance student creativity and curiosity, student satisfaction issues, various forms of assessment and metrics, engagement of underrepresented populations, and outreach drives.
- *Promotion of the societal impact of SP*: This includes privacy, ethical and security concerns, wearable devices and eHealth, global interconnections through the Internet of Things (IoT), and impact on climate change, global economy, and finance.

A coherent and comprehensive account of these issues is particularly important and timely, given the increasing exposure to popular technological advancements, such as big data, the IoT, and wearable devices. These also naturally lead to questions about the relevance of some classic subjects in modern, real-world applications.

Apart from the values specific to SP, this special issue aims to help the international community engage in education and the outreach of our discipline (including industry-run courses) to better understand, tackle, and address (through a coherent effort of international contributors) some of the key challenges the global education is facing. Indeed, the inexorable advances

Digital Object Identifier 10.1109/MSP.2021.3060277
Date of current version: 28 April 2021

in sensor technology and the IoT and the increasingly diverse forms of data acquisition have inevitably led to wider and more rapid ways in which we generate, process, and revise the notion of information. This trend is already having a major impact on how we educate and learn. Given the rich history of the SP field and the availability of competing and complementary multimedia educational resources, a common challenge in modern SP education is to produce a dynamically adjustable tradeoff, arising, as it does, from both the diversity in student learning styles and the requirements imposed by the future careers of these students.

To this end, we have identified some of the most pressing challenges the global education is facing, which include

- students communicating in a different way, which requires a rethinking of teaching practices that highlight the importance of real-time demonstrations and hands-on projects in teaching
- how to use emerging technologies to improve instruction and teaching next-generation solutions where possible (that is, educating students for jobs that currently may not even exist but will be prominent in five years or so)
- ways to bring research into the curriculum as a paradigm shift
- educating students about the importance of the completion and execution of their ideas/projects and of expressing themselves concisely and precisely through SP tools and SP ways of thinking
- the implementation of elements of service economy into electrical engineering curricula as many economies are moving away from products and into services
- enhancing awareness about the societal impact of SP education and the role of education as a key to innovation and, thus, the creation of enabling technologies for the solution of issues such as climate change, global IoT-enabled interactions, and space exploration
- the need for the reform of education, both geographically and in terms of widely accessible "global" lecture courses.

To address these challenges, we have centered this special issue of *SPM* around the following topics:

- the mitigation of issues related to the perceived difficulty of traditional SP courses, such as strategies on how to teach SP with less math and how to attract attendees from nonengineering departments
- the use of technologically orientated classrooms and emerging technologies, such as MOOCs and web resources
- metrics for success of education delivery in the after-online technology era
- using the principles of SP to improve teaching and research in related areas, such as machine learning, bioengineering, and artificial intelligence and optimization, and vice versa
- curricular changes to meet contemporary demands from industry, such as using practically relevant problems, exploring feasible extensions and new applications of the taught material, and curiosity-driven learning
- preparing students for lifelong learning and teaching lifelong fundamentals of SP and the relevance of SP with respect to technological advances
- challenges and solutions in industry-run courses—the design of short courses offered by academia for industry, government agencies, and national defense
- the role of mentorship and initiatives to encourage and motivate students in research experiences
- promoting creativity in learning, especially when applying the concepts with opportunity windows to explore entrepreneurship, possible product developments, and cross-disciplinary aspects of our work.

The timing of this special issue has been reinforced by the success of the recent special program "Celebrating Signal Processing Education" at ICASSP 2019 in Brighton, United Kingdom, which had the involvement of all of the guest editors of this special issue. This initiative has highlighted that the SP community can significantly benefit from the dissemination of ideas and practices, especially related to the recent rapid evolution of SP education. These topics are of vital importance for the future of our discipline but have not, until now, been properly addressed in a comprehensive and cohesive way in the open literature. This special issue therefore aims at providing a unifying framework to educate SP educators within the general umbrella of "Innovation Starts With Education." Before moving on to the articles in this special issue, we continue this guest editorial with a more personal "Reflections" column by two colleagues, Al Oppenheim and Tony Constantinides, who have been part of this community for more than five decades. We close with a quote from Al Oppenheim: "The role of a magician is to make simple things appear mysterious. The role of a teacher is to make mysterious things appear simple."

Guest Editors

Mónica F. Bugallo (monica.bugallo@stonybrook.edu) received her Ph.D. degree in computer science and engineering from the University of A Coruña, Spain. She is a professor of electrical and computer engineering and the vice provost for faculty affairs and for diversity, equity, and inclusion at Stony Brook University, New York, 11794, USA. She serves as the current chair of the Signal Processing Theory and Methods Technical Committee of the IEEE Signal Processing Society, senior associate editor of *IEEE Signal Processing Letters*, and associate editor of *IEEE Transactions on Signal Processing*. Her research interests include statistical signal processing, with emphasis on the theory of Monte Carlo methods and their application to different disciplines. She has also focused on science, technology, engineering, and mathematics education and has initiated successful programs to engage students at all academic stages in the excitement of engineering and research, with focus on underrepresented groups. She is a Senior Member of IEEE.

Anthony G. Constantinides (a.constantinides@imperial.ac.uk) is an emeritus professor of signal processing at Imperial College London, London, SW7 2BT, U.K. He has supervised more than 120 Ph.D. students, and his influence through his research and teaching contributions was recognized with the IEEE Kirchmayer Medal in 2012. He has further founded and served as the first president of the European Association for Signal Processing. In 1985, he was decorated by the French government with the Honour of Chevalier, Palmes Academiques and, in 1996, with the elevation to Officier, Palmes Academiques. He is a recipient of the Medal of Merit at Imperial College London. His life's work is included in the IEEE "Pioneers of Signal Processing" archives. His research interests include the interlocking areas that underpin our modern digital society of digital signal processing, image processing, graph theory, and communications. He is a Life Fellow of IEEE and a member of the Royal Academy of Engineering.

Danilo P. Mandic (d.mandic@imperial.ac.uk) received his Ph.D. degree from Imperial College London, London, SW7 2BT, U.K., in 1999, where he is currently a professor. He has served as a visiting lecturer for the "Brain Gain Programme" within World University Service and was a member of the Education Technical Committee within the IEEE Signal Processing Society. He was a recipient of the President's Award for Excellence in Postgraduate Supervision at Imperial College London in 2014. He was also a recipient of the 2019 Dennis Gabor Award, given by the International Neural Networks Society for "Outstanding Achievements in Neural Engineering," and a recipient of the 2018 Best Paper Award in *IEEE Signal Processing Magazine*. His research interests are in statistical learning and machine intelligence and their applications. He is a Fellow of IEEE and a vice president of the International Neural Networks Society.

Alan V. Oppenheim (avo@mit.edu) is with the Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts, 02139, USA. He has received multiple awards for outstanding research and teaching, including an honorary doctorate from Tel Aviv University; the IEEE Signal Processing Medal; the IEEE Education Medal; the IEEE Centennial Award; Third Millennium Medal; and the Society, Education, Technical Achievement, and Senior Awards of the IEEE Signal Processing Society. His awards at MIT for excellence in teaching, research, and mentoring include the Bose Award, the Everett Moore Baker Award, the Capers and Marion McDonald Award, the EECS Graduate Student Association Advising Award, and the MIT Freshman Advising Creative Activity Award. His research interests include signal processing and its applications. He is a member of the National Academy of Engineering and has been a Guggenheim fellow and a Sackler fellow. He is a Fellow of IEEE.

Roberto B. Togneri (roberto.togneri@uwa.edu.au) received his Ph.D. degree in 1989 from the University of Western Australia, Perth, Western Australia, 6001, Australia, where he is now an associate professor and the head of the Signal Processing and Recognition Lab. He has published more than 200 refereed papers, has been the chief investigator for three Australian Research Council Discovery Project grants since 2009, and was the area editor of columns and forums for *IEEE Signal Processing Magazine* from 1998 to 2020. His research interests include feature extraction and the enhancement of audio signals, statistical and neural network models for speech and speaker recognition, audiovisual recognition and biometrics, and language modeling and understanding. He is a Fellow of IEEE. SP

Alan V. Oppenheim and Anthony G. Constantinides

Reflections After 50-Plus Years in the Classroom

While the theme of this special issue is "Innovation Starts With Education," it is also true that "education thrives on innovation." And as technology continues to advance, new opportunities continue to present themselves for innovation in the classroom. Some remarks during the education panel at the 2019 International Conference on Acoustics, Speech, and Signal Processing started with the playful but nevertheless attention-getting comment that "we've been teaching for more than 50 years and are just realizing that, for all that time, we've been doing it wrong." That is perhaps somewhat like saying during the late 20th century that, with hindsight, when traveling from Boston to California in 1900, it would have been better to fly than to travel by rail.

The good old days and the good old ways

From ancient times (measured in decades or millennia) to the present, the passing of knowledge from "the master" to "the student" has relied on various technologies and methods. Drawing pictures in the sand and on cave walls and verbal exposition were eventually enriched by writing, the printing press, video, audio recordings, and other means by which knowledge could be stored, archived, and shared. And throughout time, the process has been dynamically augmented with experiments and real-world demonstrations to help bring concepts to life.

In more recent times, and especially before the massive disruption created by COVID-19, classroom presentation has taken a form in which the initial content exposure happens with "the master" presiding in front of a gathering of students, and, in particular, it has dealt with mathematically rigorous topics, developing in great detail—including all the epsilons and deltas and conditions for interchanging the order of integration and summation—the theorems, proofs, and examples related to the material being taught. Traditionally, in this setting, students dutifully try to copy everything, and they frequently get lost early in the presentation, confused about concepts and details, without questioning whether the math is indeed correct.

In subjects with large enrollments, lectures are generally augmented by smaller recitation sections and through even more intimate tutorial sessions and office hours with the professor. (As an interesting side comment, when one of us changed the terminology from *office hours* to *open hours* and moved the location from his office to a conference room, attendance tripled. As students pointed out, going to a professor's office can sometimes feel intimidating.) This is followed by assigned reading in textbooks and/or detailed lecture notes and homework exercises. Classroom development is typically performed with chalk on a blackboard, a marker on a white board, slides, or some combination of these. Ideally, the smaller recitation sections are highly interactive between the instructor and the students. However, in practice, too often recitation and tutorial sessions are given by relatively inexperienced graduate students who overprepare and are reluctant or unable to nimbly direct the interaction based on the needs of the students.

The use of overhead projectors and then computer-generated slides offered many opportunities to easily incorporate rich graphics and visuals (including "eye candy") into teaching. And it also provided the opportunity to focus on the highlights of mathematical derivations without "dragging" students through every small step unless there was specifically an important point to be made in doing so. Accompanying a presentation with a handout would often nicely augment a transparency or slide show and free students from having to laboriously copy everything. Unfortunately, however, instructors would often bundle the entire content of a course into static slides and then, during each semester, pull out the package without updates and without enriching it with some blackboard/whiteboard interaction. In other words, the technology had the potential to be overused, often for the convenience of the instructor. In a rapidly changing environment

Digital Object Identifier 10.1109/MSP.2020.3041242
Date of current version: 28 April 2021

 1053-5888/21©2021IEEE

such as ours, we need to consider when new ideas in teaching, motivating, and inspiring students are just substitutes for old ones and when they are another enriching dimension to be included in the tool bag.

In 1992, there was a "Reflections" column article in *IEEE Signal Processing Magazine* (*SPM*), "A Personal View of Education" [1], and another, in 2006, "One Plus One Could Equal Three (And Other Favorite Cliches)" [2], that contained some reflections on research. In preparing this article, those pieces seem as relevant today as they were then. But times have also changed in many significant ways. The profiles, expectations, and prior educational experiences of the student population are clearly different than they were two and three decades ago. And there have been significant advances in the richness of technology for sharing content, knowledge, and teachers' insights and experiences with students.

Massive open online courses and the flipped classroom

A significant step forward in incorporating technology in education was driven, in part, by the introduction of massive open online courses (MOOCs). The development of MOOCs inspired the organization of presentations into smaller modules that had video, rich graphics, concept demonstrations, and, perhaps most significantly, the automatic grading of simple exercises interspersed with the other segments. There are now many excellent MOOCs available for the signal processing community, as described, for example, in the article "MOOC Adventures in Signal Processing," published in *SPM,* in 2016 [3]. Of course, by necessity, the structure of MOOCs constrains the opportunity for rich interaction between students and teachers.

MOOCs and associated creative technologies have made an extraordinary contribution to education. They have enabled rich content presented by highly talented teachers to be accessible to anyone with Internet access anywhere in the world. And even in residential teaching environments, MOOCs have offered an opportunity for teachers who have less experience with a course's content to deepen their understanding and to incorporate elements such as the selective utilization of demonstrations, video segments, and auto-graded exercises. The development of MOOCs has also intensified the discussion of "flipped classrooms," where, before in-person sessions, students watch videos of the course content and perform simple exercises to, at a minimum, get a sense of the concepts and notation. Learners then carry out at-home or in-class applications of the methods under consideration, perhaps even employing their personal signals, e.g., their voice, electrocardiograms, and so forth (see, for example, [4]). Of course, there are many variations, from relatively strong expectations and requirements for the preclassroom components to more relaxed but encouraged assumptions.

In some ways the flipped classroom is in the spirit of the more traditional (but often ignored) suggestion to students that they spend some time with the course textbook or other reading material before coming to class. However, typical textbook material is prepared to be highly complete and detailed. Consequently, requiring textbook reading prior to any classroom exposure to context and motivation can be difficult and cumbersome and is often more meaningful after the basics have been absorbed. In any case, with whatever advance preparation students can be encouraged to do, the classroom experience becomes more than just a lecture theater: it is also a forum for inspiring, motivating, and interacting.

Well-chosen and prepared videos and autograded exercises can be enormously beneficial in acclimating students to notation and basics before a lecture or classroom interaction. The potential effectiveness of aggressive or partial classroom flipping is highly dependent on the nature of the material, the resources available to students, and the creativity and style of instructors in utilizing and building on advance preparation by the students. And again, the flipped or somewhat flipped classroom can be overdone and purposely or inadvertently take the path of being more for the convenience of the instructor than for the enhanced learning of the students.

"Necessity is the mother of invention"

As this article was being written, we were clearly experiencing another potential major step forward in incorporating technology into our teaching, precipitated by the worldwide COVID-19 crisis. During this period, schools at all levels abruptly closed their physical spaces, with the requirement to move to online platforms. This naturally meant that many of the "old" ways of delivering content—e.g., by long, detailed blackboard derivations—were, by necessity, rapidly replaced by more creative ways of presentation and engaging students. And as we all watched in real time in our respective environments, very clearly there was a lot of innovation and creative experimentation undertaken, which we all believe has impacted and will continue to influence our residential teaching methods during both the short and the long term, when life settles to whatever the new normal will be.

So, as abrupt and painful as the pandemic shutdown has been, and as extensive as the debris field will be, there are some silver linings, among them, new opportunities for presenting content and interacting with students. We've all heard the old English proverb, sometimes attributed to Plato, that "necessity is the mother of invention." With online classroom experimentation rapidly happening throughout the world, there are clearly new avenues to pursue and likely many hazards and unintended consequences. This, of course, is always the case when introducing new technology into the classroom.

Another important element in the education process is the role of mentoring, which is clearly different than that of delivering content. In this magazine, the 1992 article about education [1] emphasized the importance of live mentoring and coaching. What we are seeing at our universities during the adjustment to the pandemic are many innovative ways of having rich interaction with and among our students, evidenced, for example, by the use of "breakout rooms" (as they're referred to on the Zoom platform) and other methods of holding online open hours, maintaining accountability during exams, and so on. All these innovations have enormous potential for enhancing both residential and distance learning. Again, it's important to focus on utilizing these new resources to enhance the experience of the students rather than to benefit or provide convenience to teachers.

Textbooks

It's also important to comment on the role of textbooks. Historically, these have played an important part as a companion to the other elements of a classroom experience and as future reference material. Textbooks are often written to be highly detailed and self-contained on their own. Therefore, the content can be hard to digest during a first exposure to the material. As phrased by Andrew Wu, a Massachusetts Institute of Technology undergraduate who commented on a draft of this article:

> A problem that I personally have with textbooks is that using them can often be cumbersome. More typically, I consult a source because I have a question on some specific aspect of the material. If what I'm searching for exactly corresponds to a section in a textbook, then the textbook works well; however, if it's just a few paragraphs within a textbook, it can be tedious and cumbersome to find the exact information that I'm looking for. Modern education, to me, with its vast arrays of different technologies and methods of information delivery, offers students much more of an opportunity to learn in a more personalized way.

Of course, some of that student's concerns are mitigated by e-textbooks that support word search functions. But just as with poor search engines, it can often be difficult to find the right combination of terms to search for. Word search in an e-text is certainly a significant improvement over a poorly composed index in a hard-copy book, but it is often cumbersome and unhelpful.

While textbooks are usually not the best resource for initial exposure to material, they have always played an essential role since they can provide a more detailed exposition than is typically necessary and appropriate in the classroom. Furthermore, textbooks give students ready access to details as pupils engage material through homework exercises and related activities. Perhaps more importantly, well-written textbooks often become lifelong, trusted companions and reliable reference sources. For a host of reasons, writing and publishing textbooks—and particularly in printed, rather than electronic, form—has become less attractive to educators. Among the factors causing this are the increasingly rapid advance of the concepts, perspectives, and techniques in our field and many others and what seems like the broken "business model" of many publishers. Textbooks are typically perceived by many students as incredibly expensive to purchase, and most often they are rented or purchased used and then resold. Increasingly, there are pirated, unauthorized versions of popular texts for sale or simply posted for free on the Internet. Consequently, any financial incentives for publishers to commission, and for authors to write, textbooks are diminished.

Today's students have also grown up in an era of multitasking, for which they have developed a habit and sometimes an addiction.

Furthermore, there is an increased desirability, expectation, and requirement to incorporate hypertext links and lots of supplemental material to augment a textbook, which intensifies the overall effort on the part of authors and publishers. It currently seems unclear, at least to us, what good alternatives there are for providing students with well-written textual material while providing authors with incentives to produce it (beyond the immense satisfaction of explaining topics to a broad audience) and, indeed, for motivating publishers and publishing platforms to make content widely available at a reasonable cost. All of this again requires innovation directed toward education.

The modern student

As commented earlier, the important evolutionary changes impacting our roles as educators include the backgrounds, experiences, and expectations of students. In our own student days and throughout a major part of our personal careers, a literature search typically began with a trip to the library. Now, for all of us, including students, a literature search often starts with accessing an appropriate search engine. And in the midst of working on a problem, all of us as researchers, teachers, and students frequently find ourselves initially turning to our favorite search engine or some other online resource to direct us to the solution of, or resources related to, a problem. The vast array of online resources makes many aspects of learning and research more efficient and in many respects provides more "instant gratification." The opportunity for students to ask questions and get answers more rapidly than in decades past naturally generates a certain impatience. On the other hand, much of the information available online has not been reviewed and vetted and consequently, to some extent, it's "searcher beware."

As another aspect, students who have grown up in the era of TiVo, online and on-demand content streaming, and handheld devices have more reluctance than we did to be required to be at a specific place at a specific time. Today's students have also grown up in an era of multitasking, for which they have developed a habit and sometimes an addiction. Always having a laptop, smartphone, and smart watch nearby is wonderful for keeping up with friends, family, news, and social media, but there seems little doubt that those devices represent a strong temptation that can quickly lead students to become distracted and then lost in the classroom. This can and should be taken into account as we incorporate the rapidly expanding array of technologies for interacting with our students and delivering content.

The evolving field of signal processing

Next, we'd like to reflect on how our field has changed during the past five decades and suggest how this impacts what we choose to teach going forward. Signal processing has always been characterized by a strong symbiotic relationship between mathematics, motivating applications, and platform implementations. During the 20th century, much of the innovation was motivated by applications such as radar, sonar, avionics, communications, entertainment, and the venture into space. Platform developments were largely centered around electromagnetics, electricity, and electronics. Advances in signal processing system design and analysis relied heavily on the mathematics related to continuous functions and differential equations. And practitioners' education included, at a minimum, a firm grounding in both the fundamental mathematics and the physics related to the implementation platforms.

Toward the end of the 20th century, the digital computer moved from its role of offline analog system data analysis and simulation to a true platform for real-time and deployable signal processing systems. This opened the door to designing and implementing signal processing systems that were freed in some sense from the constraints of the physics imposed by the electronics. And, in addition to utilizing the mathematics of continuous functions, the field increasingly harnessed discrete mathematics, numerical methods, and difference equations. This transition correspondingly expanded the essential foundational mathematics, which required including and incorporating a strong understanding of linear algebra and optimization methods, statistical inference, and other approaches that are exploited in closely related fields, such as machine learning and artificial intelligence. In terms of implementing signal processing systems, computer programming skills became more central, as did a proficiency with, or at least some understanding of, integrated circuit design.

Signal processing curricula

Signals and the need for processing them arise in a very broad set of fields and disciplines, including every branch of engineering, many aspects of health science, all the physical sciences, financial data analysis, and so on. Students taking advanced undergraduate and graduate signal processing classes often have learned the prerequisites from diverse perspectives and perhaps even picked up the knowledge informally, i.e., "learning it on the street." During the first few weeks of a course, this often presents the challenge of synchronizing everyone to similar notation and perspectives. While the concepts and foundational mathematics are essentially universal across these disciplines, students will obviously relate most strongly to application contexts with which they have some familiarity.

In thinking about appropriate curricula related to signal processing, it is also important to draw a distinction between students who will be heading toward the development of signal processing tools as a technology and those who are learning signal processing primarily to apply the field's tools and methods to advance specific applications. In both cases, there is a mathematical foundation so that tools don't get misused and so that results aren't misinterpreted. (No! The MATLAB function *fft* does not generate the Fourier transform of the input signal!) Intelligent use of high-level platforms, such as MATLAB, Mathematica, and LabView, does not require an in-depth and highly sophisticated fluency with the underlying mathematics.

But interpreting results correctly does demand a basic mathematical understanding of the underlying principles as typically taught in an advanced undergraduate signals and systems course that incorporates both continuous-time and discrete-time material as well as the basic mathematics of continuous functions, linear algebra, and statistical inference. For students preparing for advanced development and research to significantly advance the technology of signal processing, an appropriate curriculum would likely also include more advanced mathematical topics, such as optimization methods, advanced statistical inference, and perhaps some functional analysis and nonlinear mathematics. For example, it is quite likely that the future of our field will involve the creative and methodical design of nonlinear systems and algorithms and the processing of signals that are best characterized on more general manifolds than Cartesian ones.

In our view, it is essential that students and practitioners advancing the technology of signal processing have real-world signals to process. Less crucial, in our opinion, is a strong commitment to advance any specific application. But it does seem indispensable that, in the process of developing creative new signal processing tools, the concepts and algorithms be tested on real as well as simulated signals. Signal models are important for developing and refining signal processing algorithms, but models are typically only approximations of real signals. It is important for students to understand the difference between signals and signal models. Anyone involved in signal processing, whether for research toward advancing the technology or for developing a specific application, needs to have real signals to process.

Some final thoughts

Our field has had a rich history, and clearly it has incredible potential going forward. There is always an opportunity for discovering or rediscovering mathematical principles that have not yet been fully exploited in the context of signal processing. And physics will continue to provide us with new ways of implementing signal processing systems. While digital platforms have played an increasingly important part in signal processing system implementation, the role of analog platforms also continues to grow, as does a mix of both. And quite likely, as the technology advances, it will become increasingly difficult to define precisely which parts of a system are considered analog and which are digital.

An additional dimension is the inevitable advancement, overlapping, and merging of multiple disciplines, offering new, rich contexts and applications on which sophisticated signal processing can have an impact. These increasing dimensions and the rapid pace of progress place further demands for the constant updating, upgrading, and modification of the material taught in classrooms. Static presentations are quickly outdated and at an accelerating pace. Industrial and societal needs are continuously evolving, pressing the need for further innovation. The confluence or divergence of different disciplines puts further pressures on the modes and content of teaching for evolving educational needs.

Authors

Alan V. Oppenheim (avo@mit.edu) received his Sc.D. degree in electrical engineering in 1964 from the Massachusetts Institute of Technology (MIT). He is a Ford Professor of Engineering and a MacVicar Faculty Fellow at MIT, Cambridge, Massachusetts, 02139, USA. He is a member of the National Academy of Engineering, Sigma Xi, and Eta Kappa Nu and the recipient of awards including the IEEE Kilby Medal, the IEEE Education Medal, the IEEE Centennial Award, the IEEE Third Millennium Medal, and the Society, Technical Achievement, and Senior Awards of the IEEE Acoustics, Speech, and Signal Processing Society. His research interests include signal processing algorithms, systems, and applications. He is a Life Fellow of IEEE.

Anthony G. Constantinides (a.constantinides@imperial.ac.uk) is an emeritus professor at Imperial College London, London, SW7 2AZ, U.K., where he is the founder and former head of the Department of Communications and Signal Processing. He is an honorary member of Eta Kappa Nu, a fellow of the Royal Academy of Engineering, and the Institution of Electrical Engineers. He has received awards including the 2012 IEEE Kirchmayer Medal, 2016 European Association for Signal Processing Meritorious Service Award, and 2018 Imperial College Medal. His research interests include the interpretation of symbols in Linear B. He is a Life Fellow of IEEE.

References

[1] A. V. Oppenheim, “A personal view of education,” *IEEE Signal Process. Mag.*, vol. 9, no. 2, pp. 69–72, 1992. doi: 10.1109/79.127285.

[2] A. V. Oppenheim, “One plus one could equal three (and other favorite cliches),” *IEEE Signal Process. Mag.*, vol. 23, no. 6, pp. 10–12, 2006. doi: 10.1109/SP-M.2006.248707.

[3] T. A. Baran, R. G. Baraniuk, A. V. Oppenheim, P. Prandoni, and M. Vetterli, “MOOC adventures in signal processing: Bringing DSP to the era of massive open online courses,” *IEEE Signal Process. Mag.*, vol. 33, no. 4, pp. 62–83, 2016. doi: 10.1109/MSP.2016.2556004.

[4] A Moniri, I Kisil, A G. Constantinides, and D P. Mandic, “Refreshing DSP courses through biopresence in the curriculum: A successful paradigm,” in *Proc. 2018 IEEE 23rd Int. Conf. Digit. Signal Process. (DSP)*, pp. 1–4. doi: 10.1109/ICDSP.2018.8631655.

SP

Keshab K. Parhi

Teaching Digital Signal Processing by Partial Flipping, Active Learning, and Visualization

Keeping students engaged with blended teaching

©SHUTTERSTOCK.COM/LUCKYSTEP

The effectiveness of teaching digital signal processing (DSP) can be enhanced by reducing lecture time devoted to theory and increasing emphasis on applications, programming aspects, visualization, and intuitive understanding. An integrated approach to teaching requires instructors to simultaneously teach theory and its applications in storage and processing of audio, speech, and biomedical signals. Student engagement can be enhanced by having students work in groups during class, where they can solve short problems and short programming assignments or take quizzes. These approaches will increase student interest and engagement in learning the subject.

Introduction

DSP is used in numerous applications, such as communications, biomedical signal analysis, health care, network theory, finance, surveillance, robotics, and feature extraction for data analysis. Learning DSP is more important than ever before because it provides the foundation for machine learning and artificial intelligence.

The DSP community has benefited tremendously from Oppenheim's views of education [1], [2] and from his many field-shaping textbooks. Teaching an engineering class in general and the DSP class in particular is very different today from 30 years ago when computers, tools, and data were not available in abundance. Shulman captures how classes with significant mathematical content were taught in the past [3]. He describes specifically how a professor teaches fluid dynamics: "He is furiously writing equations on the board, looking back over his shoulder in the direction of the students as he asks, of no one in particular, 'Are you with me?' A couple of affirmative grunts are sufficient to encourage him to continue… This is a form of teaching that engineering shares with many of the other mathematically intensive disciplines and professions; it is not the 'signature' of engineering." The author is right in that, although some instructors teach engineering this way, it is not and should not be our teaching signature. When I taught the DSP class at the University of Minnesota (UMN) using the Oppenheim–Schafer textbook [4] in the fall of 1989,

Digital Object Identifier 10.1109/MSP.2021.3052487
Date of current version: 28 April 2021

my teaching signature was close to what Shulman describes. But over the last three decades, my teaching signature has changed significantly. In this article, I will describe my teaching signature as I practice it today.

The objectives of teaching are threefold: 1) teach the necessary mathematical theory and derivations, 2) introduce sufficient applications and visualize the results by programming the application, and 3) present intuitive insight about the observations from the programming experiments. Thus, signal processing is as much about listening to sounds and visualizing temporal and spectral representations as about theoretical problem solving.

My teaching signature can be described as "blended teaching." I teach mostly by writing down the content in class. This helps the students write down what I am writing; this then helps them develop the same thought process as I do when I derive these results. I then switch to PowerPoint slides, and I show graphs and plots and play audio sounds to see how a signal sounds after a certain filtering operation or how different types of filters change the signal to different forms. We filter some music and sinusoids and discuss some MATLAB code. This blended teaching keeps the students engaged. I try to assign homework problems that relate to real applications. I remind myself of the threshold concepts [5]. While thresholds differ for different students, I try to cover as many if not all of these concepts. In signal processing, we have many threshold concepts. They range from myriads of math tricks to challenges in applying the same concepts to different problems. The most challenging part of teaching is to really pretend that we are not experts but novices. Then we can teach other novices more effectively.

Learning DSP is more important than ever before because it provides the foundation for machine learning and artificial intelligence.

Challenges in teaching DSP

When I took the DSP class at the University of Pennsylvania in 1983 using the Oppenheim–Schafer textbook [6], it was taught as an advanced graduate course at that time. The DSP class is taught today as a senior elective at most universities.

There is a desire to teach the class as a practical class, where signals, sounds, and images can be manipulated using DSP. This manipulation should be integrated into the lecture as well as homework. One of the challenges is that the textbooks are rich in theory but do not provide a sufficient number of practical applications. The textbook by Mitra, however, provides numerous applications related to multirate and sample-rate alteration [7]. Because the class is often taught with an emphasis on theory, many students lose interest in taking it. Such a class is an elective class. So to increase enrollment, students should find the class interesting and practical. We also need to train students to acquire the practical skills that will help them in their jobs in industry. However, we also need to teach the mathematical rigor for students. In the absence of an ideal textbook, this places a burden on instructors to design application examples to be covered during lectures and applications to be assigned as part of the homework. We have already seen some local success in this direction [8].

There is also a need to increase student engagement and interest. This requires instructors to deviate from traditional teaching and adopt some form of flipped teaching, where students familiarize themselves with some material before coming to the class either by reading the content or listening to video lectures [9], [10]. This frees up time in class so students can work together to solve theoretical or practical problems.

To increase attendance, short quizzes can be assigned during lectures. Assigning group quizzes can enhance student engagement by allowing students in the same group to discuss and learn from each other. Thus, taking a quiz is as much about learning as about earning a grade in the class.

Often the homework can be frustrating if the students do not learn the "tricks." Students find the lectures easy, but they find it harder to solve problems. Thus, some of the tricks to solving the problems need to be taught during lecture. This requires working out some of the problems that would have been assigned as homework. Another approach is to provide solutions to problems that are similar to the homework problems. Studying these solutions will be very helpful to the students in preparing them for their homework. The same is also true for programming problems. Starter codes for programming assignments should be provided to the students. This will help them in solving their programming assignments. Some students have strong theoretical skills but are less inclined to solve programming problems.

Finally, often there is a gap between the homework assignments and exams. Homework problems are often time consuming and require more calculations, whereas examinations cover short problems that take less time but are thought provoking and nontrivial. Students need to develop skills in solving problems that are similar to those in the tests. The aforementioned quizzes during lectures can be very helpful to students in preparing for exams.

An integrated approach to teaching

There is debate in the community about the interrelationship between innovation and education [11]. This section describes examples of how signal processing can be taught more effectively via an integrated approach that emphasizes learning of the theory, application, and intuition.

Mathematical derivations

Many decades ago, the entire class was spent on deriving the mathematical theory, and the homework problems were also mostly mathematical. Many DSP homework problems involve "tricks" that are not taught in the class, but students are expected to figure them out. As a student, I enjoyed figuring out these tricks; however, many students lose interest in learning DSP as they cannot work them out. Thus, there is a need to spend lecture time solving problems where the tricks are explained. This reduces the time required for all of the derivations. Fortunately, students can read the textbook for this part. In general, the amount of lecture time used to derive theory needs to be reduced.

The threshold concepts come into play when explaining the tricks. Students may have forgotten some of the concepts. Typically, in a junior-level class on signals and systems, I spend a week teaching functions, scaling and shifting of functions, complex numbers, and trigonometric identities.

Another aspect of deriving theory is to first explain the results intuitively and then derive the theory. At other times, it may be easier to compute the result using MATLAB and then explain the result theoretically. This achieves two objectives: students develop practical skills, and they then relate the experimental result to theory. This makes the theory more relevant. As one example, I illustrate the fast Fourier transform (FFT) properties using Table 1. Students then verify the properties using MATLAB (see Problems 5 and 6 in the section "In-Class Group Activity"). Variations of Table 1 can also be used for homework or group activity.

Applications and visualization in MATLAB

I explain several applications of the theory during the lecture. In addition, I assign programming problems for applications as part of the homework. For example, when describing digital filters in terms of, say, low-pass, bandpass, and high-pass, I take a sound or audio file, filter it with different passbands, and then listen to the filtered sound. These sounds are either embedded into the PowerPoint presentation or obtained from MATLAB during the class. I also provide the codes for the filtering operations. Students can use the code for solving their homework.

In the first DSP class, I ask students to record their speech as they describe themselves for a few minutes. They turn in their speech as part of the homework. Then, in subsequent weeks, I ask them to filter their own speech. We have also downloaded publicly available bird sounds from websites, and the students use these sounds for their homework.

In another application for speech or audio compression, we compute the FFT of the signal. Then we retain low-frequency content and compute the inverse FFT (IFFT). Then we listen to the sound. This is explored in a homework problem (Problem 1).

The MATLAB problem in Problem 1 relates to audio compression. In this problem, students explore the principles of audio compression where the high-frequency content is discarded. The MATLAB codes for the functions *fft_compress* and *fft_expand* are provided to the students.

Problem 1

1) Load the audio file, referred to as $x[n]$. Let $X[k]$ be its discrete Fourier transform (DFT). Compute $X[k]$ using the *fft* command.
2) Compress the FFT $X[k]$ using the *fft_compress* function and a percentage of compression = 10% (0.10). This retains only the first 10% of the spectrum.
3) Using the compressed sound file from step 2, apply the *fft_extract* function to reconstruct the original audio file.Save the reconstructed audio sound file and play it. Refer to this signal as $x_1[n]$. Comment on your observations.
4) Generate an error file which is the difference between the original audio file, $x[n]$, and the reconstructed audio file, $x_1[n]$. Call this error signal $e[n]$. Save the error sound file and play it. Comment on your observations.
   Comment: $e[n]$ contains the higher frequency content of $x[n]$.
5) Observe that the error signal contains frequency components in the midband and no frequency components at low frequency. Shift left the frequency components of the error signal by k_0 samples and compute the IFFT of the shifted frequency-domain signal. Save the generated sound file and play it. Call this signal $x_2[n]$. Comment on your observations.
   Comment:

$$X_2[k] = E[k + k_0]$$

$$x_2[n] = e[n]e^{-j\frac{2\pi}{N}nk_0}.$$

Table 1. The DFT properties.

$x[n]$	$X[k]$
$(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}, \mathbf{F})$	$(6a, 6f, 6e, 6d, 6c, 6b)$
(A, F, E, D, C, B)	$(\mathbf{6a}, \mathbf{6b}, \mathbf{6c}, \mathbf{6d}, \mathbf{6e}, \mathbf{6f})$
$(\mathbf{A}^*, \mathbf{B}^*, \mathbf{C}^*, \mathbf{D}^*, \mathbf{E}^*, \mathbf{F}^*)$	$(6a^*, 6b^*, 6c^*, 6d^*, 6e^*, 6f^*)$
$(\mathbf{A}, \mathbf{F}, \mathbf{E}, \mathbf{D}, \mathbf{C}, \mathbf{B})$	$(6a, 6b, 6c, 6d, 6e, 6f)$
$(\mathbf{a}^*, \mathbf{b}^*, \mathbf{c}^*, \mathbf{d}^*, \mathbf{e}^*, \mathbf{f}^*)$	$(A^*, F^*, E^*, D^*, C^*, B^*)$
$(\mathbf{a}, \mathbf{f}, \mathbf{e}, \mathbf{d}, \mathbf{c}, \mathbf{b})$	(A, F, E, D, C, B)
$(a^*, f^*, e^*, d^*, c^*, b^*)$	$(\mathbf{A}^*, \mathbf{B}^*, \mathbf{C}^*, \mathbf{D}^*, \mathbf{E}^*, \mathbf{F}^*)$
$(\mathbf{a}, -\mathbf{b}, \mathbf{c}, -\mathbf{d}, \mathbf{e}, -\mathbf{f})$	(D, E, F, A, B, C)
$(\mathbf{a}, -\mathbf{f}, \mathbf{e}, -\mathbf{d}, \mathbf{c}, -\mathbf{b})$	(D, C, B, A, F, E)
$(a, 0, b, 0, c, 0, d, 0, e, 0, f, 0)$	$(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}, \mathbf{F}, \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}, \mathbf{F})$
$(\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}, \mathbf{F}, \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{E}, \mathbf{F})$	$(12a, 0, 12f, 0, 12e, 0, 12d, 0, 12c, 0, 12b, 0)$
$(\mathbf{A}, \mathbf{0}, \mathbf{B}, \mathbf{0}, \mathbf{C}, \mathbf{0}, \mathbf{D}, \mathbf{0}, \mathbf{E}, \mathbf{0}, \mathbf{F}, \mathbf{0})$	$(6a, 6f, 6e, 6d, 6c, 6b, 6a, 6f, 6e, 6d, 6c, 6b)$
$(A/6, 0, F/6, 0, E/6, 0, D/6, 0, C/6, 0, B/6, 0)$	$(\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}, \mathbf{f}, \mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}, \mathbf{f})$
$(A/12, F/12, E/12, D/12, C/12, B/12, A/12, F/12, E/12, D/12, C/12, B/12)$	$(\mathbf{a}, \mathbf{0}, \mathbf{b}, \mathbf{0}, \mathbf{c}, \mathbf{0}, \mathbf{d}, \mathbf{0}, \mathbf{e}, \mathbf{0}, \mathbf{f}, \mathbf{0})$
$(\mathbf{D}, \mathbf{0}, \mathbf{E}, \mathbf{0}, \mathbf{F}, \mathbf{0}, \mathbf{A}, \mathbf{0}, \mathbf{B}, \mathbf{0}, \mathbf{C}, \mathbf{0})$	$(6a, -6f, 6e, -6d, 6c, -6b, 6a, -6f, 6e, -6d, 6c, -6b)$

Example given in class: Let the discrete Fourier transform (DFT) of a six-point complex sequence (a, b, c, d, e, f) be another complex sequence (A, B,C, D, E, F). Then complete the table below. The sequences in bold were given, and students were asked to find the corresponding pairs; the sequences in red correspond to solutions.

The signal $x_2[n]$ is complex and differs from $e[n]$ and is a modulated version of $e[n]$. Thus, if we listen to its magnitude, it will sound different from $e[n]$.

6) Multiply the signal $x_2[n]$ with a complex exponential $(e^{j(2\pi/N)k_0 n})$, where k_0 corresponds to the shift in frequency performed in step 5. Save the generated sound file and play it. Call this signal $x_3[n]$. Comment on your observations.
Comment:

$$x_3[n] = x_2[n]e^{j\frac{2\pi}{N}nk_0} = e[n].$$

$x_3[n]$ is the same as $e[n]$.

Solution: Solutions with MATLAB codes are provided in the supplementary materials that appear with this article on IEEE *Xplore*. ■

I introduce practical applications while describing theoretical concepts. For example, while introducing the definition of *autocorrelation* of a real signal, I provide examples of photoplethysmogram (PPG) and respiration-rate signals. Then I discuss how to compute the heart rate and respiration rate from these two signals using autocorrelation by looking at the zero-crossings. We then compute the DFT of the signals and verify if the frequency obtained by the autocorrelation is the same as that from the DFT. This is illustrated in Problem 2, where the PPG signal is used to compute the heart rate. The respiration-rate signal is not included in Problem 2, but the approach is similar. This helps connect the theoretical expression for autocorrelation to a practical application.

Problem 2

The PPG signal captured from a sensor at a 100-Hz sampling frequency is provided in this problem. The length of the signal is 1,024 samples (10.24 s). The data file (ppg_100hz_1024samples.csv) for this problem is given in the supplementary materials that appear with this article on IEEE *Xplore*. The data are part of a PPG signal (ppg_100hz_1024samples.csv). The PPG signal is used in this problem to compute the heart rate.

1) *Compute the 1,024-point FFT of the signal and plot the absolute values of the single-sided FFT with a stem plot*: Find the frequency in hertz of the highest magnitude in the FFT of the PPG. Note that the frequency corresponding to the highest magnitude represents the heart rate.
2) *In this part, we compute the heart rate using autocorrelation*: This is accomplished by finding the difference of the first and third zero-crossings, which corresponds to the time period of the signal. This information is used to compute the heart rate in hertz from the PPG.

Solution: The MATLAB code for this problem is available in the supplementary materials that appear with this article on IEEE *Xplore*. A diagram containing results from the two parts of the problem is presented in Figure 1. The heart rate can be estimated by:

1) *Finding the highest peak from the DFT spectrum*: The fundamental frequency is 1.0742 Hz for the PPG signal, which results in a heart rate of 64.45 beats per minute (bpm).
2) *Considering the interval between first and third zero-crossings*: A lag difference of 94 at a 100-Hz sampling rate = 0.94 s or 63.8 bpm. Note that the values from steps 1 and 2 are almost the same. ■

Fortunately, numerous large collections of data and signals are now publicly available. These data and signals can be used as part of the homework or class projects. In my DSP class, I have used intracranial electroencephalogram signals for seizure detection from the UPenn and Mayo Clinic's Seizure Detection Challenge on Kaggle [12]. Students use the same signals for solving different

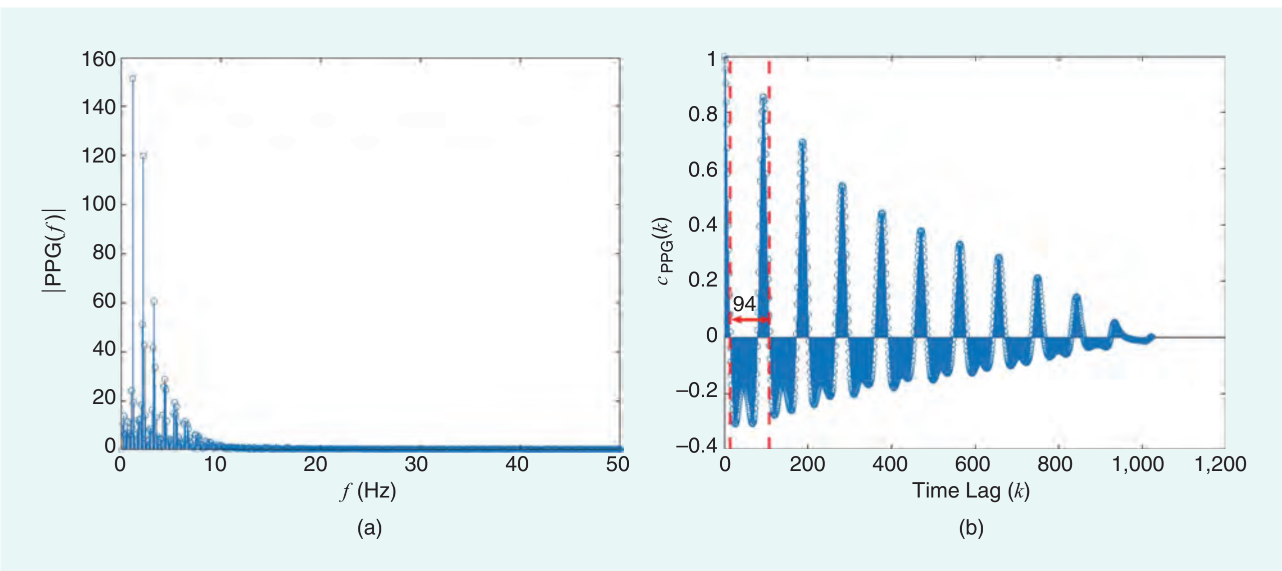

FIGURE 1. (a) The DFT and (b) normalized autocorrelation for the PPG signal.

programming problems assigned over many weeks and compute time-domain and frequency-domain features. These problems are described in "Programming Assignments Using Intracranial Electroencephalogram Data" (see Problems S1–S3).

The roles of theory and application sometimes can be interchanged. We first describe an application using MATLAB. Then we make an observation and then derive the theory.

Intuitive insight

It is important to explain the results from MATLAB experiments intuitively. For example, the effects of scaling and shifting a signal in the spectral domain can be quickly observed. The theory can then be explained. I included the following Problem 3 as part of a homework assignment.

Problem 3

Two signal processing systems are shown in Figures 2 and 3, where $x_1[n]$ and $x_2[n]$ are audio sounds, and $H(z)$ is a 100th-order finite-impulse response low-pass filter with cutoff frequency $\pi/2$. For each system, load the input audio sounds, use MATLAB to obtain the sounds $z[n]$, $y_1[n]$, and $y_2[n]$, and listen to these sounds. Use the *freqz* command to plot the spectrum of the sounds $x_1[n]$, $x_2[n]$, $z[n]$, $y_1[n]$, and $y_2[n]$ for each system. Compare the output signals obtained using the two DSP systems.

Solution: Multiplication of $x_2[n]$ by $(-1)^n$ results in a shift in the frequency domain by π. Thus, the signal $z[n]$ contains the audible $x_1[n]$ along with the shifted version of $x_2[n]$, which is inaudible to the human ear. The second multiplication of $z[n]$ by $(-1)^n$ results again in a shift in the frequency domain by π, thus making it possible to listen to $x_2[n]$ at $y_2[n]$. However, $y_2[n]$ also contains the high-frequency content of $x_1[n]$. This can be avoided by band-limiting the input signals using system 2 shown in Figure 3. ■

The intrigue of the previous problem lies in the fact that the sound $z[n]$ does not seem to contain $x_2[n]$, whereas it is audible in $y_2[n]$. Students think that the signal $x_2[n]$ is lost, and they are surprised that it can be recovered. I then explain this mathematically and intuitively. This problem can illustrate the basic concepts of audio steganography.

Blended teaching and active learning

Almost all students today have their own laptops that they can bring to class. Thus, it is easy for them to learn in an active

Programming Assignments Using Intracranial Electroencephalogram Data

This is a description of programming assignments that use intracranial electroencephalogram (EEG) data from the Seizure Detection Challenge on Kaggle organized by the University of Pennsylvania and the Mayo Clinic [12]. The students were assigned one specific subject and a specific electrode/channel signal from that subject. The students were asked to extract various time- and frequency-domain features and comment on the suitability of these features for discriminating ictal (during seizure) and interictal (baseline) clips to detect seizures. The sample solutions for these problems, provided in the supplementary materials that appear with this article on IEEE *Xplore*, use the training data from channel 1 for the EEG clips from Dog 1, which has a total of 596 clips of which 178 are ictal, i.e., those corresponding to seizures. Each clip is a 1-s recording.

Problem S1. This problem explores time-domain signal processing. Extract and plot the eight time-domain features listed below for the assigned subjects. Use the *stem ()* command to plot the value of each feature for all of the clips. Observe the plots and comment on whether the given feature could be used to detect seizures.

1) *Measures of central tendency*: arithmetic mean, median, and mode.
2) *Energy*: the energy for a sequence $x[n]$ of length N is given by

$$E = \sum_{n=1}^{N} |x[n]|^2.$$

3) *Total length of the curve (sum of distances between successive points)*:

$$L = \sum_{n=2}^{N} |x[n] - x[n-1]|.$$

4) *Hjorth parameters*:

$$\text{Activity}(x(t)) = \text{var}\{x(t)\}$$

$$\text{Mobility}(x(t)) = \sqrt{\frac{\text{var}\left\{\frac{d}{dt}x(t)\right\}}{\text{var}\{x(t)\}}}$$

$$\text{Complexity}(x(t)) = \frac{\text{Mobility}\left(\frac{d}{dt}x(t)\right)}{\text{Mobility}(x(t))}.$$

Solution: The MATLAB codes for this problem are provided in the supplementary materials that appear with this article on IEEE *Xplore*. The results varied based on the assigned subject and electrode. In many cases, the energy, total length, and Hjorth activity were good indicators of seizure. ■

Problem S2. In this problem, we will generate an analytic signal from the EEG signal. The analytic signal is a complex signal whose real part is the signal itself and whose imaginary part is the signal filtered by the Hilbert transform filter. Using this, we generate two new features listed below. Plot the features using the *stem ()* command learning environment where they can write short programs or solve short problems during the class. Students can learn from each other by working in groups.

Flipped classes have been used to teach DSP effectively [13]. At UMN, I taught the undergraduate DSP class EE-4541 in an active learning classroom in the fall of 2013 [14], [15]. The students sat around tables in groups of three. The classroom was equipped with a camera for the instructor and there were TV screens near each table, as shown in Figure 4. I sometimes used pen and paper to derive or explain theoretical results. At other times, I used PowerPoint slides. Students had access to my PowerPoint slides a week before the class, and they were asked to review them before coming to class.

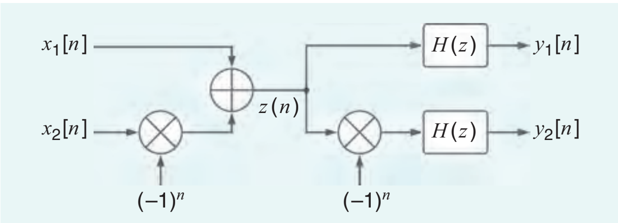

FIGURE 2. Signal processing system 1 for Problem 3.

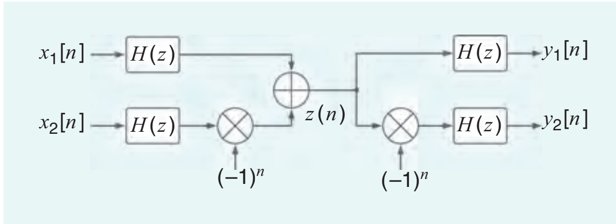

FIGURE 3. Signal processing system 2 for Problem 3.

FIGURE 4. An active learning classroom.

and comment on whether the given feature could be used to detect seizures.

1) *Mean instantaneous amplitude of delta band*: Filter the original EEG clips using a bandpass filter to obtain the signal in the delta band (1–4 Hz). Using the Hilbert transform, obtain the discrete-time analytic signal (complex valued), the magnitude of which provides the instantaneous amplitude. Use its average as the feature.
2) *Mean instantaneous frequency of alpha band*: Filter the original EEG clips using a bandpass filter to obtain the alpha-band (8–12 Hz) signal. Using the Hilbert transform, obtain the discrete-time analytic signal (complex valued), the angle of which provides the instantaneous phase. The derivative of the unwrapped instantaneous phase scaled by the sampling frequency yields the instantaneous frequency. Use its average as the feature.

Solution: The MATLAB codes for this problem are provided in the supplementary materials that appear with this article on IEEE *Xplore*. The results varied based on the assigned subject and electrode. In many cases, the mean instantaneous amplitude of the delta band is a good indicator of seizure. ■

Problem S3. We explore the following three methods listed to observe the power spectral density (PSD) of the EEG clips and find out if PSD is a useful feature for seizure detection.

1) *Spectrogram*: Combine all the clips (ictal followed by interictal) to form a single time series. Use the spectrogram command with a window of 100 sample segments and an 80-sample overlap to view the frequency spectrum. Show the output as a surface plot with time on the x-axis, frequency on the y-axis and spectrum (in decibels) along the third axis.
2) *Welch PSD estimate*: Combine the ictal clips and interictal clips separately to form two different time series. Using the *pwelch* command, obtain and plot the PSD estimate of the two signals on the same graph for a normalized frequency range of $[0-\pi]$. Use the smoothdata() function to smooth the plot and identify the normalized frequency range that has the maximum difference between the PSD estimates of the ictal and interictal series.
3) *Average PSD as a feature*: Using the *stem ()* command, plot the average of the PSD estimate obtained using the *pwelch* command for each clip. Observe and comment on whether this feature could be used to detect seizures.

Solution: The MATLAB codes for this problem are provided in the supplementary materials that appear with this article on IEEE *Xplore*. The results varied based on the assigned subject and electrode. In many cases, the PSD estimate is a good indicator of seizure. ■

I assigned a quiz in the first class. The students were divided into three groups based on their performance in the quiz: top, middle, and bottom. Groups consisting of three students each were created by randomly picking one student from each of the top, middle, and bottom groups. This course was taught twice per week, where each class was of 75-min duration. Out of the 75 min, the last 15 min were reserved for either a group activity or a group quiz. During the group activity, the students were assigned short problems and short MATLAB assignments to work on in their groups. In this approach, a student who needed help could learn from another student in the group. The group quizzes also consisted of short problems and short MATLAB programming problems. The group quiz and group activity alternated from one class to another during the semester. At the end of the group activity, I was able to provide intuitive insights and solutions to the problems at the end of the lecture.

In-class group activity

I designed the group activity problems such that students could first either learn the tricks needed to solve problems or compute the final result by MATLAB before the theory was presented. Other problems were designed to use MATLAB to verify what was learned from theory. Some examples of group activity are described next.

Problem 4

Consider the following sinc function:

$$x[n] = \left(\frac{\sin\frac{n\pi}{4}}{n\pi}\right).$$

Using MATLAB, plot the discrete-time Fourier transform (DTFT) of the 10 signals listed below.

1) Plot DTFT $X(e^{j\omega})$ using MATLAB.
2) Let $x_1[n] = x[n-10]$. Plot $X_1(e^{j\omega})$.
3) Let $x_2[n] = x[-n]$. Plot $X_2(e^{j\omega})$.
4) Let $x_3[n] = nx[n]$. Plot $X_3(e^{j\omega})$.
5) Let $x_4[n] = e^{j(n\pi/6)}x[n]$. Plot $X_4(e^{j\omega})$.
6) Let $x_5[n] = (-1)^n x[n]$. Plot $X_5(e^{j\omega})$.
7) Let $x_6[n] = x[n] * x[n]$. Plot $X_6(e^{j\omega})$.
8) Let $x_7[n] = x^2[n]$. Plot $X_7(e^{j\omega})$.
9) Let $x_8[n] = x[2n]$. Plot $X_8(e^{j\omega})$.
10) Let $x_9[n] = \begin{cases} x[n/2], & n \text{ is even} \\ 0, & n \text{ is odd} \end{cases}$.

Solution: The MATLAB codes to this problem are provided in the supplementary materials that appear with this article on IEEE *Xplore*. The students were asked to explain the discrepancy between impulse magnitudes in the MATLAB result and the theoretical result for step 4 ($201/2\pi$ versus 1). ■

As another example of group activity, I will ask the students to observe the properties of the DFT by solving Problem 5 using MATLAB. Then they solve Problem 6 based on their observations from Problem 5. Once the students understand the properties, I derive some of them theoretically in class. Finally, students explore the use of the DFT properties shown in Table 1 as part of their homework. In this table, the sequences in red correspond to the solutions and are assigned as homework.

Problem 5

Evaluate the following using MATLAB:

1) FFT [2, 3, 4, 5, 6]
2) FFT [FFT[2, 3, 4, 5, 6]]
3) FFT [2, 3, 4, 5, 6, 0, 0, 0, 0, 0]
4) FFT [2, 3, 4, 5, 6, 2, 3, 4, 5, 6]
5) FFT [2, 2, 3, 3, 4, 4, 5, 5, 6, 6]
6) FFT [4, 5, 6, 2, 3]
7) FFT [2, − 3, 4, − 5, 6].

Solution: The FFT of a sequence can be computed using the *fft* command in MATLAB. ■

Problem 6

Let $(a, b, c, d, e) \Leftrightarrow (A, B, C, D, E)$. Write general expressions for the FFTs listed below in terms of (A, B, C, D, E) based on the observations from Problem 5.

1) FFT [FFT[a, b, c, d, e]]
   Solution: [$5a, 5e, 5d, 5c, 5b$]. ■
2) FFT [a, b, c, d, e, 0, 0, 0, 0, 0]
   Solution: [$A, *, B, *, C, *, D, *, E, *$]. Here $*$ denotes interpolated values and hence cannot be expressed in terms of A, B, C, D, E. ■
3) FFT [$a, a, b, b, c, c, d, d, e, e$]
   Solution: [A, B, C, D, E, A, B, C, D, E]$[1 + e^{-j(k\pi/5)}]$ for $k = 0, 1, \ldots, 9$. ■
4) FFT [$a, b, c, d, e, a, b, c, d, e$]
   Solution: [$2A, 0, 2B, 0, 2C, 0, 2D, 0, 2E, 0$]. ■
5) FFT [d, e, a, b, c]
   Solution: $[A, Be^{-j(4\pi/5)}, Ce^{-j(8\pi/5)}, De^{-j(12\pi/5)}, Ee^{-j(16\pi/5)}]$ using the property $x([n-2])_{<5>} \leftrightarrow e^{-j(4k\pi/5)}X[k]$. ■
6) FFT [$a, -b, c, -d, e$]
   Solution: Interpolation with two-and-a-half-sample delay. Here the input can be expressed as $(-1)^n x[n] = x[n]e^{-j(2\pi n/5)(5/2)}$. ■
7) FFT [a, e, d, c, b]
   Solution: $[A^*, B^*, C^*, D^*, E^*]$, using the property $x([-n])_{<5>} \leftrightarrow X^*[k]$. ■

In-class group quiz

Students take a group quiz lasting 15 min once a week. This engages the students in the group to solve the problems together. It also reduces the pressure of taking quizzes for individual students. An example of a group quiz is given in Problem 7. Group quizzes help the students to prepare for the examinations.

Problem 7

Evaluate the following. Note that these time-domain convolution problems are easier to solve in the frequency domain.

1) $\dfrac{\sin\left(\frac{n\pi}{4}\right)}{n\pi} * \dfrac{\sin\left(\frac{n\pi}{8}\right)}{n\pi}$.

Solution: The frequency-domain representation of a sinc signal is a rectangular function. The solution then involves multiplying two rectangular functions and then taking an inverse Fourier transform, which is another sinc function. [See Figure 5(a).] ■

2) $$\frac{\sin\left(\frac{n\pi}{4}\right)}{n\pi} * \left(\frac{\sin\left(\frac{n\pi}{2}\right)}{n\pi} - \frac{\sin\left(\frac{n\pi}{3}\right)}{n\pi}\right).$$

Solution: See Figure 5(b). ■

3) $$\frac{\sin\left(\frac{n\pi}{4}\right)}{n\pi} * \left(\delta[n] - \frac{\sin\left(\frac{n\pi}{8}\right)}{n\pi}\right).$$

Solution: The same as before; we take advantage of the fact that the Fourier transform of the δ function is 1. See Figure 5(c). ■

4) $$\left(\delta[n] - \frac{\sin\left(\frac{n\pi}{4}\right)}{n\pi}\right) * \left((-1)^n \frac{\sin\left(\frac{n\pi}{4}\right)}{n\pi}\right).$$

Solution: Note, $e^{-j\pi n} = (-1)^n$. Thus, the signal is shifted in the frequency domain by π. See Figure 5(d). ■

5) $$\left((-1)^n \frac{\sin\left(\frac{2n\pi}{3}\right)}{n\pi}\right) * \sin\left(\frac{n\pi}{4}\right).$$

Solution: See Figure 5(e). ■

Evaluation metrics

Group comparison

Data were collected from EE-4541 students at UMN in Fall 2012 (41 respondents out of 77 enrolled) and Fall 2013 (62 respondents out of 87 enrolled). The data collection was approved by the Institutional Review Board (IRB) at UMN under the "Exempt" category. The metrics for Fall 2012 serve as the baseline for the comparison.

The two groups of students did not differ significantly on any available demographic variables, including undergraduate–graduate status, year in the university, ethnicity, sex, age, cumulative grade point average, and composite ACT score. We can conclude that, as far as can be determined from the available data, the students in the two sections of EE-4541 can be validly compared to one another.

Outcome analyses

The metrics used to understand the efficacy include *engagement, enrichment, flexibility, effective use, classroom/course fit, confidence*, and *student learning outcome* (*SLO*). For each measure, a number of criteria and questions were chosen for the students, and they were asked to grade each question as Strongly Agree, Agree, Disagree, or Strongly Disagree, corresponding to numerical scores of 3, 2, 1, and 0, respectively. A brief description of each of the metrics is presented next.

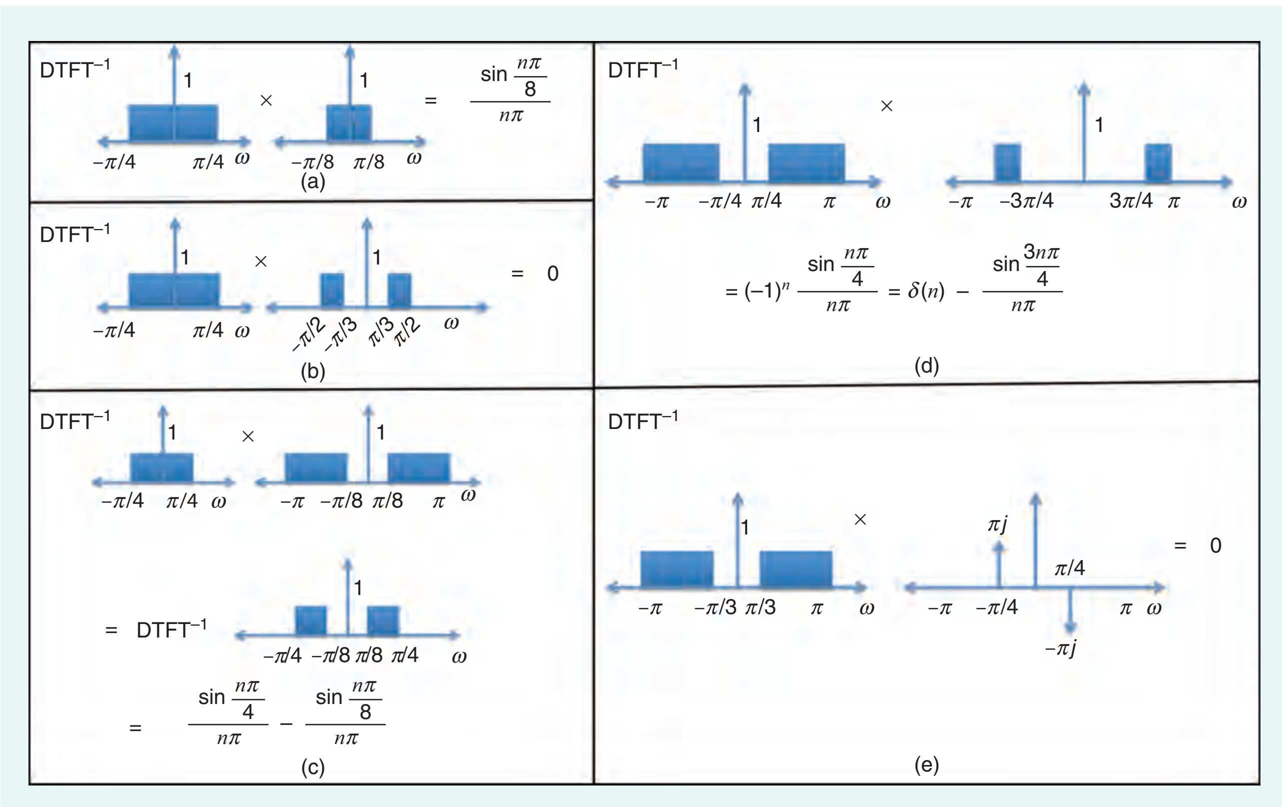

FIGURE 5. The Problem 7 solutions.

1) *Engagement*:
- Encourages my active participation
- Promotes discussion
- Helps me develop connections with my classmates
- Helps me develop connections with my instructor
- Engages me in the learning process.

2) *Enrichment*:
- Enriches my learning experience
- Makes me want to attend class regularly
- Increases my excitement to learn.

Flexibility:
- Facilitates multiple types of learning activities
- Nurtures a variety of learning styles.

Effective use:
- The instructor is effective in using the technology available in the classroom for instructional purposes.
- The instructor is effective in using the classroom for instructional purposes.

Classroom/course fit:

The criteria used for the classroom/course fit metric are listed below:
- The classroom is an appropriate space in which to hold this particular course.
- The in-class exercises for this course are enhanced by the features of this classroom.

Confidence:

The students rate the course based on the following criteria:
- Course helps develop confidence in working in small groups
- Helps students develop confidence in analyzing
- Helps student develop confidence in presenting
- Helps develop confidence in writing
- Improves confidence that the student can speak clearly and effectively.

3) *SLO*:
- Helps me develop professional skills that can be transferred to the real world
- Helps me to define issues or challenges and identify possible solutions
- Prepares me to implement a solution to an issue or challenge
- Helps me to examine how others gather and interpret data and assess the soundness of their conclusions
- Deepens my understanding of a specific field of study
- Assists me in understanding someone else's views by imagining how an issue looks from his or her perspective
- Helps me to grow comfortable working with people from other cultures
- Improves my confidence that I can speak clearly and effectively
- Encourages me to create or generate new ideas, products, or ways of understanding
- Prompts me to incorporate ideas or concepts from different courses when completing assignments
- Enabled the instructor to make intentional connections between theory and practice in this course.

Table 2. Student evaluation metrics.

Variable	Semester	*N*	Mean Score	*p*
	Fall 2012	41	2.345	0
Engage	Fall 2013	62	2.960	
	Fall 2012	41	2.565	0.019
Enrich	Fall 2013	62	2.854	
	Fall 2012	41	2.329	0
Flexibility	Fall 2013	62	3.169	
	Fall 2012	41	2.793	0.028
Effective	Fall 2013	62	3.129	
	Fall 2012	41	2.598	0.004
Fit	Fall 2013	62	3.024	
	Fall 2012	41	2.327	0.001
Confidence	Fall 2013	62	2.672	
	Fall 2012	41	2.497	0.007
SLO	Fall 2013	62	2.768	

Bivariate tests

Independent-samples t-tests were conducted to compare the learning metrics of the students taught in the Fall 2013 semester using partial flipping and active learning versus those in the Fall 2012 class with traditional learning. The results are summarized in Table 2. On all aggregated variables derived from student responses, statistically significant differences (at the $p < 0.05$ level or better) were found between the mean scores of the two groups, favoring the Fall 2013 class (see Table 2). The group-level difference was the highest in the categories of engagement, flexibility, and confidence. The next highest categories include classroom/course fit and SLO. There is still room to improve the scores in the enrichment and effective use categories.

Current trends

While MATLAB is used in many universities and industries, it is not an open source environment. There is great interest in teaching DSP using Python or Octave as they do not require licenses. However, students have to write code from scratch for many DSP functions, unlike in MATLAB, where students can use numerous in-built functions. Nevertheless, using Python is more desirable as most open source libraries for machine learning functions are written in Python. There is also growing interest in teaching DSP for embedded systems such as smartphones so that students can design apps for cell phones [16]. For example, they can write DSP programs in Python for smartphones to analyze biomedical signals such as electrocardiograms. Most commercial products like smartwatches already have this capability. We should create DSP lab courses to teach app design for either Android or iOS operating systems to prepare students for the rapidly changing job environment.

The entire world was disrupted by COVID-19 during the initial white paper submission of this article (2020

February) and submission of the full paper (submitted June 2020 and revised 2020 December). Almost all classes in the second half of the Spring and Fall semesters of 2020 were taught using remote learning. This provided a challenge and opportunity to redesign various courses. Many laboratories were redesigned so that the students could perform the experiments at home. I provided lecture notes and recorded videos of the EE-4541 class from Fall 2017 [10] to my students in the Fall 2020 semester. All programming problems discussed in this article were assigned as homework in the Fall 2020 class. Active learning is still possible using breakout rooms in a remote learning environment such as Zoom; however, it is better suited for an in-person class.

Conclusions
We argue that DSP can be taught effectively by using visualization, active learning, and partial flipping. Application examples enable visualization, where instructors can play different sounds and illustrate plots of time-domain and spectral-domain features. This will increase student engagement, interest in the class, and understanding of the subject. Use of speech, audio, and biomedical signals in the class and as part of the homework can connect the theory to applications and better prepare students for jobs in industry. Future DSP textbooks should include application examples and connect the theory to applications. However, instructors can use the application examples presented in this article to supplement the textbook.

Acknowledgment
The partial-flipping experiment was carried out when the author participated in the 2012–2013 Office of Information Technology (OIT) Faculty Fellowship Program at UMN. The EE-4541 class was taught in the Fall 2012 semester without flipping (baseline) and in Fall 2013 with partial flipping in an active learning classroom. The author is most grateful to Kim Wilcox and Lauren Marsh from UMN OIT, who coordinated the program. The data collection for the OIT Faculty Fellowship Program was approved by the UMN IRB (principal investigators: D. Christopher Brooks, Kim Wilcox). The statistical analysis of the metrics was carried out by J.D. Walker. The evaluation reported in the section “Evaluation Metrics” is his contribution to this article. All faculty fellows met multiple times during the two-year program. The faculty group discussed many aspects of teaching, ranging from use of apps to applying concepts of game theory [17] and crowdsourcing in teaching. The author is grateful to Bhaskar Sen for his help in preparing this article. Zisheng Zhang and Sai Sanjay Balaji were teaching assistants for the class and developed the MATLAB codes. The author thanks Prof. M. Sabarimalai Manikandan of the Indian Institute of Technology, Bhubaneswar, India, for providing the PPG application example and the PPG signal. This article has supplementary downloadable material available at 10.1109/MSP.2021.3052487, provided by the author.

Author

Keshab K. Parhi (parhi@umn.edu) received his Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, in 1988. He has been with the University of Minnesota, Minneapolis, Minnesota, 55455, USA, since 1988, where he is currently the Distinguished McKnight University Professor and Edgar F. Johnson Professor of electronic communication in the Department of Electrical and Computer Engineering. He has published more than 650 papers, is the inventor of 31 patents, and has authored the textbook *VLSI Digital Signal Processing Systems* (Wiley, 1999). He served as the editor-in-chief of *IEEE Transactions on Circuits and Systems, Part I* during 2004 and 2005. He is a Fellow of IEEE and a fellow of the American Association for the Advancement of Science, the Association for Computing Machinery, and the National Academy of Inventors.

References

[1] A. V. Oppenheim, “A personal view of education,” *IEEE Signal Process. Mag.*, vol. 9, no. 2, pp. 69–72, 1992. doi: 10.1109/79.127285.

[2] A. V. Oppenheim, “One plus one could equal three (and other favorite cliches),” *IEEE Signal Process. Mag.*, vol. 23, no. 6, pp. 10–12, 2006. doi: 10.1109/SP-M.2006.248707.

[3] L. S. Shulman, “Signature pedagogies in the professions,” *Daedalus*, vol. 134, no. 3, pp. 52–59, 2005. doi: 10.1162/0011526054622015.

[4] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing*, 1st ed., Englewood Cliffs, NJ: Prentice Hall, 1989.

[5] J. Meyer and R. Land, *Threshold Concepts and Troublesome Knowledge: Linkages to Ways of Thinking and Practising Within the Disciplines.* Edinburgh: Univ. of Edinburgh, 2003.

[6] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*, 1st ed., Englewood Cliffs, NJ: Prentice Hall, 1975.

[7] S. K. Mitra, *Digital Signal Processing: A Computer-Based Approach*, 4th ed., New York: McGraw-Hill, 2010.

[8] S. Kanna, W. von Rosenberg, V. Goverdovsky, A. Constantinides, and D. P. Mandic, “Bringing wearable sensors into the classroom: A participatory approach [SP education],” *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 110–130, 2018. doi: 10.1109/MSP.2018.2806418.

[9] T. A. Baran, R. G. Baraniuk, A. V. Oppenheim, P. Prandoni, and M. Vetterli, “MOOC adventures in signal processing: Bringing DSP to the era of massive open online courses,” *IEEE Signal Process. Mag.*, vol. 33, no. 4, pp. 62–83, 2016. doi: 10.1109/MSP.2016.2556004.

[10] K. K. Parhi. UMN EE-4541 Digital Signal Processing (Fall 2017). Accessed: Dec. 5, 2020. [Online Video]. Available: https://www.youtube.com/playlist?list=PLT1QAv48lhQLPwgL4Y–XTzKeQBViAYKv

[11] V. Solo, M. S. Greco, D. Mandic, P. Djuric, A. Spanias, and M. F. Bugallo, “Innovation starts with education: ICASSP 2019 education panel [SP forum],” *IEEE Signal Process. Mag.*, vol. 36, no. 5, pp. 135–147, 2019. doi: 10.1109/MSP.2019.2923280.

[12] “UPenn and Mayo clinic’s seizure detection challenge.” Kaggle. https://www.kaggle.com/c/seizure-detection (accessed Sept. 12, 2020).

[13] B. Van Veen, “Flipping signal-processing instruction [SP education],” *IEEE Signal Process. Mag.*, vol. 30, no. 6, pp. 145–150, 2013. doi: 10.1109/MSP.2013.2276650.

[14] K. K. Parhi. “Teaching digital signal processing by partial flipping.” Cultivating Change Series. https://cultivatingchangeseries.com/ffp/teaching-digital-signal-processing-by-partial-flipping/ (accessed Feb. 1, 2020).

[15] K. K. Parhi. Teaching Digital Signal Processing using Blended Teaching. (Jan. 29, 2014) Accessed Feb. 1, 2020. [Online Video]. Available: https://youtu.be/gwrISY4OVe8

[16] R. W. Schafer. “EE 264 digital signal processing.” Stanford. https://web.stanford.edu/class/ee264/ (accessed Feb. 1, 2020).

[17] R. Epper, A. Derryberry, and S. Jackson. “Game-based learning: Developing an institutional strategy https://library.educause.edu/resources/2012/8/gamebased-learning-developing-an-institutional-strategy (accessed: June 14, 2020).

SP

Sally A. Male, Roberto B. Togneri, and Lu C. Jin

Novice to Postgraduate Researcher Perceptions of Threshold Concepts and Capabilities in Signal Processing

Understanding students' and researchers' perspectives

©SHUTTERSTOCK.COM/LUCKYSTEP

Signal processing is an engineering discipline known to involve abstract and complex concepts. Curriculum development should be informed by an understanding of the most critical and challenging learning in the field. Threshold concept theory and threshold capability theory provide a framework describing the features of the most critical and challenging learning in any discipline. The framework describes the effort of overcoming thresholds as troublesome, with a process that is often messy and long. Five coursework master's students, six postgraduate research students, and five academics were interviewed about their experiences with threshold concepts in signal processing. Two major threshold concepts were identified: time–frequency transformation and discretization. Self-regulated learning through years was needed to overcome the thresholds. Based on students' comments, the following are recommended to support learning in signal processing: integrated units, an introduction to how signals can be represented and why signal processing is used, examples of real applications, visualizations, practical laboratory exercises with prework, small applied projects throughout units, ample sample problems, the development of learning communities through consistent class groups, and opportunities to ask questions. Coursework and research students reported developing efficacy in self-directed learning as a consequence of overcoming threshold learning in signal processing.

Introduction

Signal processing is an engineering discipline involving significant conceptual hurdles, such as Fourier transforms. Additionally, signal processing has widely diverse applications (e.g., noise removal, source localization, telecommunications, adaptive filtering, image enhancement, and source coding). Students often struggle with these challenges. Recently, educators redesigned curricula by using at least one of two strategies. The first is flipping, meaning that students access material before going to class and engage in interactive, facilitated learning during class [1], [2]. The second approach is to enhance students' engagement through practice [3].

Our previous investigation into improving teaching and learning in signals and systems [4] involved simple strategies to better engage students, avoid stretching pupils too far at once, and provide

Digital Object Identifier 10.1109/MSP.2021.3055201
Date of current version: 28 April 2021

a more inclusive learning environment. We engaged student peer facilitators with first-hand knowledge from their own recent experiences of being introduced to signals and systems. Additionally, the third author taught a communications unit in intensive mode, meaning that the course lasted eight weeks instead of a full semester, increasing the continuity of engagement [5]. Any such efforts to improve teaching and learning in signal processing should be informed by compelling evidence of the most critical and challenging concepts and capabilities for students to master, how acquiring this knowledge is troublesome, and how pupils surmount the challenges.

Previous studies considered the concepts students need to learn in signal processing. Martinez-Torres et al. [6] used multivariate statistical analysis to map concepts based on the experience and knowledge of teachers. Wage et al. [7] developed a signal processing concept inventory standardized test that measured all the important constructs in signal processing. Identifying every essential notion is not sufficient. It is also important to identify the most critical and challenging aspects of learning.

Theoretical framework: Threshold concept and threshold capability theory

Especially when developing interactive learning opportunities, it is necessary to identify the concepts and capabilities that are the most transformative and critical to future knowledge and practice in the field. In threshold concept theory, these ideas and proficiencies are known as *threshold concepts* and *threshold capabilities* [8], [9]. Since they are transformative, threshold concepts and capabilities, also known as *thresholds*, are usually troublesome for students. The term *threshold* is not used in the sense of being a level, as is often the meaning in engineering. Instead, it refers to a gateway to future learning and practice. In mastering threshold concepts and capabilities, students grow to be comfortable applying them. By identifying thresholds and the features that make them difficult, educators can design curricula to focus class time on the most essential and taxing learning [10], [11].

The theoretical framework of threshold concepts and capabilities has been shown to be valuable for curriculum development. In 2012, a team at the authors' university, including the first and second authors, informed the design of new foundation engineering units for students in the first years of all engineering disciplines at the school [12]. Researchers and curriculum developers seek to identify thresholds in a discipline, and they explore how students find thresholds troublesome and how learners overcome those obstacles. A popular method for identifying thresholds is through interviews and focus groups with students and teachers. Threshold concept theory describes common characteristics that can be used to help identify thresholds [13]. Among other common features, the learning associated with overcoming a threshold is usually "irreversible" [14, p. 110], often connecting previously isolated concepts for the learner and frequently enhancing the learner's use of language. The transformative nature of thresholds means that overcoming a threshold fosters capabilities that the learner did not have before.

Especially relevant to how students overcome thresholds, the theory refers to the "liminal space" [12, p. 398] as the state that a student is in when he or she has become aware of a concept or capability but not yet become comfortable with it. Studying students' transition through the liminal space can inform curriculum development. The framework explains the importance of considering students' understanding before they entered the liminal space. Because threshold concepts and capabilities are difficult, it can take more than one unit (also known as a *course* or a *paper*) for students to become comfortable with the information and skills. Therefore, educators provide opportunities for students to revisit thresholds during a degree program. However, some students require even more time and overcome thresholds only after they have been working for an extended period.

Previously identified thresholds in signal processing

Studies of thresholds generally involve students and teaching team members as participants and often focus on a single unit. In their study of how to support students learning electrical engineering, Carstenson and Bernhard, in 2008 [15], analyzed video recordings of laboratory classes. They explored the links that students made, including investigating the questions that students asked and the connections that learners made between physical and various models of systems. Carstenson and Bernhard coined the term *complex concept* [15, pp. 146–147], which links multiple ideas. They concluded that, in addition to designing the sequence of concepts to be taught, it is necessary to craft curricula such that students inevitably undertake learning activities that force them to make connections between subject areas. In a foundation electronics unit, Harlow et al., in 2011 [16], used student surveys, student focus groups, interviews with teaching team members, and an analysis of assessment responses and grades to identify threshold concepts. Harlow and her coresearchers found that the lecturer was not easily able to identify the troublesome aspects of learning for the students because the concepts were not troublesome for the instructor. Consistent with Carstenson and Bernhard, the challenges identified by Harlow et al. were related to connections between models of systems.

In 2012, the first and second authors, along with others, identified foundation engineering threshold concepts and capabilities through a two-step process involving students and educators participating in interviews and/or focus groups, followed by the negotiation of identified thresholds [12], [17], [18]. Findings related to signal processing were consistent with previous studies. "Abstraction, modeling, and theories," depending on other thresholds, such as "model evaluation," "visual representation of concepts and systems," "describing systems mathematically," and "relating mathematical representations of systems to physical systems," were identified as thresholds [19, pp. 20–28]. In 2017, Reeping et al. [20] held focus groups with electrical and computer engineering faculty members to identify big ideas that should be in the curriculum and how those areas relate to threshold concepts identified in the literature. In terms of signal processing, they recognized "frequency domain," "complex analysis" (linked to Fourier analysis), and "demodulation" [20, p 5].

The research outlined in the preceding paragraphs identified threshold concepts related to signal processing. However, the participants were limited to teachers and students. As noted, academics can have difficulty identifying thresholds because they have become familiar with them through experience [15], [16]. Conversely, students are thought to have difficulty identifying thresholds because they cannot yet determine what they have not understood [15]. It is known in threshold concept theory that students can take longer than one unit, and even beyond graduation, before they become comfortable with threshold concepts. Therefore, it is reasonable to expect that the perspectives of postgraduate practitioners would complement the views of students and academics. However, although practitioners are often involved in curriculum design, they have been included in threshold concept studies only in limited disciplines, not including signal processing. An exploration of the variation in perspectives of threshold concepts from coursework students to practitioners was expected to enhance the understanding of the most critical and challenging learning in signal processing and how students' development progresses.

Research questions

To inform signal processing curriculum development, we investigated perceptions of students at various levels, from coursework to research, and academics as practitioners who apply signal processing, addressing the following questions:

1) What are the thresholds (concepts and/or capabilities) in signal processing experienced by students of various levels, teachers, and practitioners?
2) How do coursework students', graduate research students', and academics' perspectives of thresholds in signal processing vary with levels of experience?

Method

Semistructured interviews were conducted with coursework students, postgraduate research students, and academics who were teaching signal processing or supervising pupils who were using signal processing in research. Interviews are used in educational research to deeply explore topics, especially those that are not well understood. In semistructured interviews, a researcher poses planned questions and expands on them with additional prompts, even following unexpected topics that emerge and are of interest. The purpose is to reach a thorough understanding of the views of a selected group of participants. In this study, the questions focused on participants' experiences of learning about signal processing and, where relevant, observations of students they had taught or supervised. By asking about observations of others' learning, the researchers investigated the experiences of many more people than participated in the study.

Participant recruitment

Research designs involving interviews do not necessarily need large numbers of subjects. This study used a purposive sample with intentional diversity of experience among the participants. Consistent with ethics approval, invitations to participate were emailed to students enrolled in relevant units, peer tutors in these units, laboratory demonstrators, academics teaching the units, doctoral students identified by their supervisors as using signal processing, and academics identified as using signal processing.

Interviews

Interviews were conducted one-to-one with the exception of one focus group involving three coursework students. The first author conducted all interviews except one in which the participant was a current student of all three authors. A researcher with electronic engineering and threshold concept research expertise conducted that interview. The interviewer explained threshold concepts and capabilities as transformative and asked the participants to identify knowledge or skills they had experienced and/or observed students to experience. For each identified threshold, the participant was asked to outline the concept or capability and describe the following:

- the transformation provided by the threshold
- how the threshold was troublesome
- what he or she or others do or did before overcoming the threshold
- any concepts that the threshold connected
- any expanded use of language associated with the concept
- how he or she overcame or observed others overcoming the threshold
- whether he or she needed to revisit the threshold
- how the threshold was used
- barriers and useful approaches to overcoming the threshold.

Interviews were recorded and transcribed.

Participants

Participants included five coursework master's engineering students, six postgraduate research students, and five academics. The five students had various coursework experience of signal processing. Of three relevant units—signals and systems, signal processing, and communication systems—one student was undertaking the first, one student had completed the first and was undertaking the second, one student had completed the first two and tutored in the first, one student had completed and demonstrated in the second (having completed units equivalent to the first elsewhere), and one student had completed all three units. Three of the coursework students were male, and two were female.

The postgraduate research students were pursuing doctoral degrees, applying signal processing in diverse fields, including acoustic signal processing, optical coherence tomography, gravity wave detection, and communications. Two had been involved in teaching signal processing. Four were male, and two were female. The academics were using signal processing in their research, supervising coursework students from various disciplines of engineering and/or physics who used signal processing in their research projects, and supervising postgraduate research students who used signal processing. Three of the academics taught units related to signal processing. Four were male, and one was female, and two were the second and third authors of this article.

Analysis

The transcripts were analyzed in four stages. First, the potential thresholds identified in each interview were documented. Second, those areas were validated against the compulsory criterion for a threshold concept or capability, namely, being transformative. Third, the validated thresholds were rationalized to combine similar ones and remove repetition. Fourth, the structure of the relationships between thresholds and the variation in thresholds across the levels of experience among participants was mapped.

Findings and discussion

The interviews revealed different thresholds challenging students at various stages. Students and academics reported the need to have confidence with the mathematical concepts of integrals, logarithms, complex exponentials, random processes, power spectral densities, covariance, and correlations. In the coursework stage, students reported a threshold related to understanding why signal processing is necessary and the different forms in which signals could be generated and represented. Working with time–frequency transforms, in particular, the Fourier transform, was the most discussed threshold capability. Potential thresholds nested within this were visualization, complex exponentials, and modeling. Discretization was the second-most commonly discussed threshold. Nested within this were convolution, sampling, and windowing. Two postgraduate students and one coursework student identified program coding (e.g., MATLAB) as a threshold capability that transformed how they developed signal processing. One academic identified model fitting as a threshold. Phase and decibels were identified as threshold concepts by different academics. Participants' reports reliably revealed that some threshold concepts in signal processing are especially troublesome. Even highly capable students, who became teachers, required years after graduation before they were comfortable with the most challenging threshold concepts in signal processing.

Two major identified thresholds

Two threshold concepts received significantly more attention in the interviews than any of the others, as described in the following.

Threshold 1: Time–frequency transformation

The time–frequency transformation threshold concept was identified most generically and concisely by academic 2:

> There's definitely a threshold concept behind being able to see data in multiple domains. And the Fourier transform, the time-to-frequency swap in the Fourier domain, is one of those.

Transformative

The concept is transformative in the sense that it enables the learner to complete tasks that were not previously possible, as described by academic 5:

> With understanding it, it opens up a lot of techniques that we use to process, particularly, images or data. So, if you can take some time series information, like … the power of the light coming in as a function of time, and say you were just interested in only the green lines and how that changes with a function of time, the Fourier transform would enable you to cut away the blue and the red and just leave you with the green.

Participants reported that the concept enabled them to carry out design in signal processing. Graduate student 3 said,

> You know the transformation first, and then you can design the signal processing system by yourself.

The transformation is clear in a quote from graduate student 5:

> It was pretty much a … watershed moment, I think, when it sort of clicked, and it came from pretty much a year of hard slog, working long hours … the normal sort of eight-to-five thing in the Ph.D., plus time at night and on the weekend, self-learning and everything … to get comfortable with how the optical coherence tomography (OCT) technique worked. And there was sort of a huge moment when I really I taught myself … Fourier analysis. And once I got it into my own terms, and I understood it, it was this huge moment where I began to understand how OCT worked. What was once this enormous beast just simplified everything, and then it also simplified all of the stuff I'd struggled with … the very first time I even picked up [units], when I first had to deal with … transforms or Fourier analysis, in which I was so confused. And so … I didn't … have this sort of threshold learning until I was already well and truly into a Ph.D.

Troublesome features

The transformation concept is troublesome because it is abstract, as described by graduate student 5:

> I think the hardest part is being able to—I don't actually, I can't think of signals in terms of frequency. It doesn't make any sense. It's not. It's a lot. If you think of a signal varying as a function of time you can size it .… It goes up or down or whatever .… We're used to thinking of things in a progression of time. Whereas thinking of things in terms of a frequency is pretty weird and doesn't come naturally to me anyway.

Of the issue, academic 1 said:

> You just have to understand that [the complex Fourier transform] is an abstraction. So, things like negative frequency, things like why are things complex, also the symmetry properties …. If you can understand you're going from a real world to an abstraction … of the physical system to be able to design and analyze it more efficiently .… Complex Fourier transforms is a clear case where you do that.

Participants also identified visualizing the physical systems in each domain as an additional threshold, and this could be considered a threshold that increases the troublesome nature of the Fourier transform.

Threshold 2: Discretization

The second major threshold concept was the consequence of discretizing a process. Processing discrete rather than continuous data introduces possibilities for error. Academic 2 described what happens when students do not understand the concept:

> [Students] might understand the pure mathematical Fourier transform. And they know how to invoke the fast Fourier transform button in MATLAB, but they have absolutely no practical grasp on all the things that go wrong when suddenly the thing you are processing is a string of numbers rather than an analytical function.

Transformative

Academic 2 further identified what can be done with an understanding of the discretization threshold concept:

It helps you understand the impact of various processing steps of a signal: why one processing step might lose information, why it is you need a certain bandwidth or a certain something else to capture that information, or you might get hurt.

An example of getting hurt is aliasing caused by an incorrect application or implementation of sampling.

Troublesome features

Academic 1 reported,

> Students don't realize [that] when you go to the discrete domain, you're not talking about time, you're talking about samples of data.

Academic 2 explained:

> What's happening is you are, in order to compute each one point down here, you're taking an integral from minus infinity, the first infinity of the first function multiplied by the second function .… Engineers usually think of it as little dips and bits and hats. And they go, "Why is the radio not working?" Because you have, you've failed to capture … the hat out here.

Thresholds nested within the concept of discretization include convolution, windowing, and sampling. Convolution involves combining two signals in the time domain. It is difficult to visualize the process. Aliasing occurs if there is an error in how the signal was sampled. Academic 2 described windowing:

> When you discretize data, you necessarily have a finite chunk of it. And yet, the Fourier transform is defined as an integral across all time. What happens when you have this finite chunk of data (anything you've necessarily measured is finite), when you take a Fourier transform, there is an implicit infinite replication of this thing. And sometimes that will cause a horrible outcome. So, you go, "Well, gee, I'm going to have to want to squish the horrible outcome. And I do it by premultiplying this by something, which is a window .…" We say the data has been windowed, meaning it is finite, but I change the shape of the window by attenuating the guys on the end, perhaps, or adding huge numbers of zeros or some other technique to try and stop one of the bad things that happened because of the implicit infinite replication.

How students overcame the thresholds

Deep, self-regulated learning

Students described taking responsibility for learning about signal processing. Many participants reported that overcoming the thresholds took extensive reading and checking. A need in the workplace and teaching signal processing to others were catalysts in overcoming thresholds.

Barriers to learning

Students learn about transforms in first-year math, and the signals and systems unit is in the third year. This was noted as a problem by an academic and by a coursework student, who said:

> There's a massive gap between … when you first start Laplace transforms and all these transforms in mathematics … and then when you do the signal processing signal and system unit. So, I think most people have already forgotten that. And I think, from the get-go, they've already struggled to remember what they were back then. So maybe I think it would be better if we did … the refresher unit or a refresher week, when you do repeat the whole transformations again.

Transition through the liminal space

The transition through the liminal space toward becoming comfortable with threshold concepts in signal processing is long and difficult. The interviews revealed that even successful research students and academics take years to understand them. Participants reported remaining unclear about some elements of the two identified threshold concepts. Overcoming the thresholds involved deep, self-regulated learning approaches through years. The study of threshold concepts by Harlow et al. found that the teacher was not aware of the difficulties in understanding faced by students because the instructor did not have similar struggles. In contrast, academic 4 reported taking five to 10 years and applying significant effort to understand the frequency domain and convolution:

> So, I, back at my, the way I was taught, and I didn't, I didn't get it. I never understood, not frequency domain or convolution, after leaving university. I, it probably took me five years, 10 years before I was really comfortable … with exactly what it meant, and that, for me, came out of really deconstructing a lot of the math and realizing that all of the math .… I taught the unit in [University X] … and I think I must have pulled out about six different books on signal processing.

Overcoming the thresholds in signal processing was associated with a transformation in students' confidence in their ability to learn. One graduate described himself as not learning quickly, and another described himself as not a logical thinker. Those participants were not concerned about this at the time they were interviewed, and instead reported that it assisted them in explaining concepts to others. The interviews of these students and a coursework student reveal that the learners developed self-efficacy through developing understanding and skills in signal processing. Graduate student 2 said:

> Well, now, I'm doing pretty fundamental work in signal processing in our field, and I've got publications, so I'm at a point where … I can do it .… It's not like: he's the one that might be able to. It's like: I can; I've got the publications to back it up .… So, I had the ability all along, but I just didn't get it back when I was an undergrad. I just didn't understand it.

Graduate student 4 remarked:

> When I was an undergrad, I relied on people to teach me things. I was relying on the course notes, the tutorial examples, the tutorial questions, and the lecturer. And the lecturer was going too fast. The pace was too fast for me. And I didn't understand the concepts well enough for the examples to help me, and I couldn't synthesize the information out of the lecture notes enough, and whilst I probably could have if I had the time to, I probably could have figured out myself, there were … two things stopping me from doing that. The

> first thing was time commitments with other units. So, I was doing three other units at the time .… That occupies a lot of space in your mind and also time studying for tests and exams and everything; you can't do that. And also, I hadn't really got to the point where I was able to teach myself things.

Students reported the following approaches to overcoming the thresholds in signal processing: reading books and accessing online resources, including videos; using visual representations; solving practice problems and finding opportunities to ask questions of teachers; learning about applications; teaching one another in a close cohort; being challenged early (through projects during a unit and by preparing before laboratory sessions); and teaching.

Visual representations
On the topic of visual representations, coursework student 1 said:

> Having a physical, like, graph of a signal, say, and then manipulating with math and then producing a different graph, it just tends to help me remember how, how each of this, the filters are supposed to work, and then how … I'm supposed to implement them later.

Learning community
Graduate student 6 reflected on learning in a cohort, remarking:

> I had a class of possibly 30-ish total, or it was, maybe it was two classes of 30 .… It meant that we stuck in groups, and we worked in groups, and we all taught each other .… You know, the lecture did that. Then we taught each other, and we did that.

Learning by teaching
Graduate student 6 described learning by teaching:

> And each time I did explain it, I seemed to learn a little bit more about it and go, "Oh, well, that happens because of this." The next time, you know, the students would have different problems .… I guess I learned from the mistakes of myself and others, and as a teacher, I get to see a lot more mistakes.

Being challenged early
Students need to enter the liminal space, in which they tackle the troublesome elements of a threshold concept. If they reach the troublesome elements only after class, they have lost the opportunity to resolve the trouble with others. Coursework student 1 described the need to start some preparation before laboratory sessions. For similar reasons, in another comment, the student suggested multiple small projects during semester instead of one project at the end, saying:

> I find … I need to … go through the content of the lab and try and do, like, at least a quarter of it before I go into the lab, and honestly just to have some understanding of what I'm supposed to be doing .… You can … read the manual and things like that like, but I find that, in signal processing, that I can't just … research from the unit reader; like, the content that I'm supposed to know … I always feel like out of my depth .… That's maybe an assignment that you do on MATLAB; say that's not in the, that's not in the 3-hour timed environment .… There are prelab activities, but it's mostly just, it's mostly, like, solving this probability .… It's just a little too basic in order to really properly help us.

Variation across participants
The coursework students identified fundamental aspects of the first major threshold concept of time–frequency transformation, such as what frequency is, what a signal is, how a signal can be represented, and why signals are processed. The graduate students and the academics focused heavily on conceptual understanding and alternatives they had observed to be held by the coursework students. They were able to identify students' actions before and after the pupils reached understanding, and they were able to recall what they had done to help. The inclusion of graduate students who applied signal processing in diverse applications was intentional. The study found that the threshold concepts in signal processing were similar across applications. This is reassuring for curriculum design.

Comparison with previous studies
The first transformation threshold concept is consistent with topics identified by Reeping et al. [20] and Carstenson and Bernhardt [15]. Those studies, with their separate methods, revealed a different understanding of the concepts. This study has demonstrated the long, troublesome nature of the experience of developing an understanding of the concepts. Similarly, the comments from participants in this study are consistent with the identification by Harlow et al. [16] of threshold concepts in electrical and electronic engineering associated with making connections between physical systems and models.

Recommendations
Most importantly, by consulting coursework students and graduate research students, this study discovered that it takes learners longer than one unit and even longer than a coursework degree program to cross the liminal space for major threshold concepts in signal processing. This means that units should be designed together to support students in continuing to develop knowledge across multiple units. It also means that significant time should be allocated to allowing students to focus on the major threshold concepts. If units are cluttered with too many topics, students are likely to complete degree programs without achieving the most critical learning.

The study also revealed practices likely to improve students' transition through the liminal space. A thorough introduction is necessary, covering frequency, what signals are, how signals can be represented, and why signals are processed. Students should have ready access to examples of applications, visual representations, and practice problems. Educators should design units such that students reach the challenging aspects of concepts sufficiently early in a semester, not at the end. Students should be supported in developing a learning community [21] to enhance their interactions and opportunities by teaching one another. This could be done by limiting students' movement between labs and/or by designing activities to support student interactions. Educators should design connections between units. It would be helpful if students connected what they learn in mathematics with applications in signal processing. If possible, units should be completed close in time. The revision of mathematics within units would help.

Limitations and future research

Studies with 10 participants are not designed to reach generalizable results. The sample was sufficient to be confident about the major threshold concepts because they were reliably raised by participants. However, the interview transcripts were rich in potential threshold concepts and, as noted, some concepts were identified by one person only. Therefore, additional interviews and further analysis may reveal more threshold concepts. The study used interviews. Methods involving observations may discover topics that the participants did not have the awareness to self-report.

Conclusions

The approach was overdue to inform curriculum design that takes account of signal processing students' development, including the experiences of graduates who use signal processing. This study contributes to understanding how students' perspectives of threshold concepts and capabilities develop.

Acknowledgments

The authors gratefully acknowledge the student participants, the academic participant who chose to remain anonymous, Fiona Panther, and Jonathan Scott for contributing to this research and for conducting interviews.

Authors

Sally A. Male (sally.male@uwa.edu.au) received her Ph.D. degree in engineering education in 2011 from the University of Western Australia, Perth, Western Australia, 6001, Australia, where she holds the chair in engineering education. She is editor-in-chief of *Australasian Journal of Engineering Education* and associate editor of *Journal of Engineering Education*. She is a fellow of Engineers Australia and the recipient of a 2019 ICASSP Education Innovation Award and the 2017 Engineers Australia Medal. Her research interests include engineering practice and capabilities, work-integrated learning, curriculum development, and gender inclusion in engineering education.

Roberto B. Togneri (roberto.togneri@uwa.edu.au) received his Ph.D. degree in 1989 from the University of Western Australia, Perth, Western Australia, 6001, Australia, which he joined in 1988 and where he is now an associate professor and the head of the Signal Processing and Recognition Lab. He has published more than 200 refereed papers, been the chief investigator for three Australian Research Council Discovery Project grants since 2009, He was the area editor, columns and forums, for *IEEE Signal Processing Magazine* from 1998 to 2020. His research interests include feature extraction and the enhancement of audio signals, statistical and neural network models for speech and speaker recognition, audiovisual recognition and biometrics, and language modeling and understanding. He is a Fellow of IEEE.

Lu C. Jin (lu.jin@uwa.edu.au) received his Ph.D. degree in electronics engineering from the University of Western Australia, Perth, Western Australia, 6001, Australia, in 2012, where he is now a senior lecturer. His research interests include signal processing and underwater wireless communications. He is a Member of IEEE.

References

[1] B. V. Veen, "Flipping signal-processing instruction [SP Education]," *IEEE Signal Process. Mag.*, vol. 30, no. 6, pp. 145–150, 2013.

[2] W. U. Bajwa, "On 'flipping' a large signal processing class [SP education]," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 158–170, 2017. doi: 10.1109/MSP.2017.2698074.

[3] T. Schäck, M. Muma, and A. M. Zoubir, "Signal processing projects at Technische Universität Darmstadt: How to engage undergraduate students in signal processing practice," *IEEE Signal Process. Mag.*, vol. 34, no. 1, pp. 16–30, 2017. doi: 10.1109/MSP.2016.2619218.

[4] R. Togneri and S. A. Male, "Signals and systems: Casting it as an action-adventure rather than a horror genre," in *Proc. 44th Int. Conf. Acoustics, Speech Signal Process.*, Brighton, U.K., 2019, pp. 7859–7863.

[5] D. Huang and L. Jin, "Communication technologies course with intensive mode teaching: A case study," in *Proc. IEEE Global Eng. Educ. Conf.*, Dubai, UAE, 2019, pp. 133–138.

[6] M. R. Martinez-Torres, F. J. B. Garcia, S. L. T. Marin, and S. G. Vazquez, "A digital signal processing teaching methodology using concept-mapping techniques," *IEEE Trans. Edu.*, vol. 48, no. 3, pp. 422–429, 2005. doi: 10.1109/TE.2005.849737.

[7] K. Wage, J. R. Buck, C. H. G. Wright, and T. B. Welch, "The signals and systems concept inventory," *IEEE Trans. Edu.*, vol. 48, no. 3, pp. 448–461, 2005. doi: 10.1109/TE.2005.849746.

[8] J. H. F. Meyer and R. Land, "Occasional report 4," Enhancing Teaching-Learning Environments in Undergraduate Courses, May 2003. [Online]. Available: http://www.etl.tla.ed.ac.uk/docs/ETLreport4.pdf

[9] C. Baillie, J. A. Bowden, and J. H. F. Meyer, "Threshold capabilities: Threshold concepts and knowledge capability linked through variation theory," *Higher Educ.*, vol. 65, no. 2, pp. 227–246, 2013. doi: 10.1007/s10734-012-9540-5.

[10] D. Perkins, "Constructivism and troublesome knowledge," " in *Overcoming Barriers to Student Understanding: Threshold concepts and troublesome knowledge*, J. H. F. Meyer and R. Land, Eds. London and New York: Routledge, 2006, pp. 33–47.

[11] G. Cousin, "Neither teacher-centred nor student-centred: Threshold concepts and research partnerships," *J. Learning Develop. Higher Educ.*, vol. 1, no. 2, pp. 1–9, Feb. 2010.

[12] S. A. Male and C. A. Baillie, "Research guided teaching practices: Engineering thresholds; an approach to curriculum renewal," " in *Cambridge Handbook of Engineering Education Research*, A. Johri and B. Olds, Eds. Cambridge, U.K.: Cambridge Univ. Press, 2014, pp. 393–408.

[13] S. A. Male, "Guide for engineering educators on curriculum renewal using threshold concepts," *Australian Government Office for Learning and Teaching*, Sydney, 2012. [Online]. https://ltr.edu.au/resources/PP10_1607_Baillie_Inventory_2012.pdf

[14] C. Zander, J. Boustedt, A. Eckerdal, R. McCartney, J. E. Mostrom, M. Ratcliffe, and, and K. Sanders, "Threshold concepts in computer science: A multi-national empirical investigation," in *Threshold Concepts within the Disciplines Some Examples from Electrical Engineering*, R. Land, J. H. F. Meyer, and J. Smith, Eds. Rotterdam, The Netherlands: Sense, 2008, pp. 105–118.

[15] A.-K. Carstenson and J. Bernhard, "Threshold concepts and keys to the portal of understanding," in *Threshold Concepts within the Disciplines Some Examples from Electrical Engineering*, R. Land, J. H. F. Meyer, and J. Smith, Eds. Rotterdam, The Netherlands: Sense, 2008, pp. 143–154.

[16] A. Harlow, J. Scott, M. Peter, and B. Cowie, "'Getting stuck' in analogue electronics: Threshold concepts as an explanatory model," *European J. Eng. Educ.*, vol. 36, no. 5, pp. 435–447, 2011. doi: 10.1080/03043797.2011.606500.

[17] K. M. Quinlan, S. A. Male, C. A. Baillie, A. Stamboulis, J. Fill, and Z. Jaffer, "Methodological challenges in researching threshold concepts: A comparative analysis of three projects," *Higher Educ.*, vol. 66, no. 5, pp. 585–601, Nov. 2013. doi: 10.1007/s10734-013-9623-y.

[18] S. A. Male and C. A. Baillie, "Threshold capabilities: An emerging methodology to locate curricula thresholds," in *Proc. Res. Eng. Educ. Symp.*, Madrid, 2011, pp. 907–915.

[19] S. A. Male, "Integrated engineering foundation threshold concept inventory," *Australian Government Office for Learning and Teaching*, Sydney, 2012, Available: https://ltr.edu.au/resources/PP10_1607_Baillie_Threshold_concept_2012.pdf

[20] D. Reeping, L. D. McNair, M. Wisnioski, A. Y. Patrick, T. L. Martin, L. Lester, B. Knapp, and S. Harrison, "Using threshold concepts to restructure an electrical and computer engineering curriculum: Troublesome knowledge in expected outcomes," in *Proc. IEEE Front. Educ. Conf.*, 2017, pp. 1–9.

[21] E. Wenger, *Communities of Practice – Learning, Meaning and Identity*. New York: Cambridge Univ. Press, 1998.

SP

Setareh Maghsudi, Andrew Lan, Jie Xu,
and Mihaela van der Schaar

Personalized Education in the Artificial Intelligence Era

What to expect next

©SHUTTERSTOCK.COM/LUCKYSTEP

The objective of personalized learning is to design an effective knowledge acquisition track that matches the learner's strengths and bypasses his/her weaknesses to ultimately meet his/her desired goal. This concept emerged several years ago and is being adopted by a rapidly growing number of educational institutions around the globe. In recent years, the rise of artificial intelligence (AI) and machine learning (ML), together with advances in big data analysis, has introduced novel perspectives that enhance personalized education in numerous ways. By taking advantage of AI/ML methods, the educational platform precisely acquires the student's characteristics. This is done, in part, by observing past experiences as well as analyzing the available big data through exploring the learners' features and similarities. It can, for example, recommend the most appropriate content among numerous accessible ones, advise a well-designed long-term curriculum, and connect appropriate learners by suggestion, accurate performance evaluation, and so forth. Still, several aspects of AI-based personalized education remain unexplored. These include, among others, compensating for the adverse effects of the absence of peers, creating and maintaining motivations for learning, increasing the diversity, removing the biases induced by data and algorithms, and so on. In this article, while providing a brief review of state-of-the-art research, we investigate the challenges of AI/ML-based personalized education and discuss potential solutions.

Introduction

The last decade has witnessed an explosion in the number of web-based learning systems due to increasing demand in higher-level education, the limited number of teaching personnel, advances in information technology and AI, and, more recently, COVID-19. In the past few years, to enhance conventional classrooms, to bridge the constraints of time and distance, and to improve fairness by making high-quality education accessible, most universities have integrated massive open online course (MOOC) platforms in their education systems. Also, several schools have added online labs to their structures, where

Digital Object Identifier 10.1109/MSP.2021.3055032
Date of current version: 28 April 2021

students, especially those who cannot access physical labs, can perform experiments.

In addition, there has been significant growth in the development of other online educational tools that simplify learning. These include, for example, the software used for text summarization in different domains and also to produce questions and tests, followed by evaluation, which can be of great assistance to not only students but also to teachers. Several advantages of these systems over traditional classroom teaching are that 1) they provide flexibility to the student in choosing what and when to learn, 2) they do not require the presence of an interactive human teacher, and 3) often, the capacity in terms of the number of participants is significantly larger than the synchronous-presence teaching form. Figure 1 shows the baseline ecosystem of online personalized education, including all the stakeholders, together with the crucial factors and performance metrics.

However, currently available online teaching platforms have significant limitations. To a large extent, personalized education has been limited mainly to a specific type of recommender system, although its potential goes far beyond advising a series of lectures on an online platform that might be interesting to a specific user. One fundamental difference between existing recommender systems and personalized education is the optimization objective: The former focuses on some form of user engagement to maximize profit, which is system centric and relatively easy to quantify, whereas the latter focuses on some form of learning outcomes, which is student centric and hard to define.

ML/AI-enabled education is a response with great potential to overcome the current shortcomings. It creates a new and more flexible learning technology genre that adapts to student learning and allocates resources as obliged. It takes advantage of the strengths of both online tools and individual tutoring. As such, AI-enabled personalized education promises to yield many of the benefits of one-on-one instruction at a per-student cost similar to that of large university lecture classes. The system applies to both online courses and courses with a hybrid of classroom and online instruction. As displayed in Figure 2, ML/AI-enabled education comprises a large set of decision-making strategies that collectively map the available data together with the individual features to a variety of personalized educational materials and recommendations.

Data can be collected on performance in both traditional assignments (problem sets, computer programs, and laboratories) as well as online exercises and tests. It includes built-in assessment tools as an essential part of its optimization of lesson sequences. As such, it supports the educational community in developing new teaching modalities in a broad range of disciplines. However, despite intensive research efforts conducted during this decade, a variety of aspects of personalized

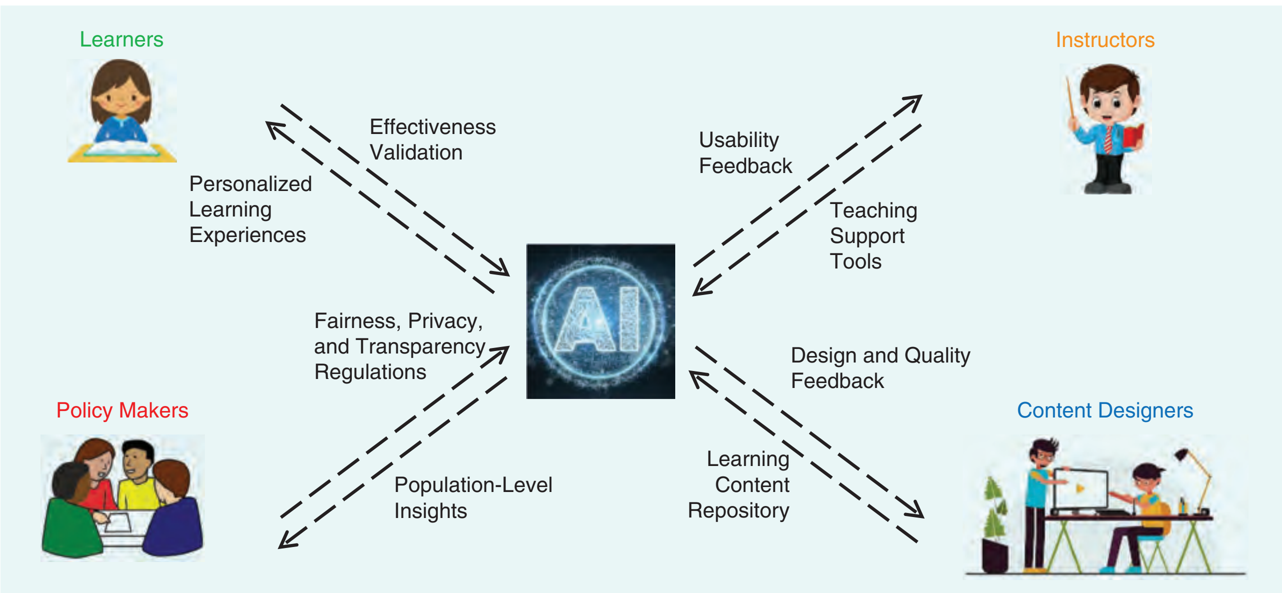

FIGURE 1. The baseline ecosystem of AI-empowered personalized education.

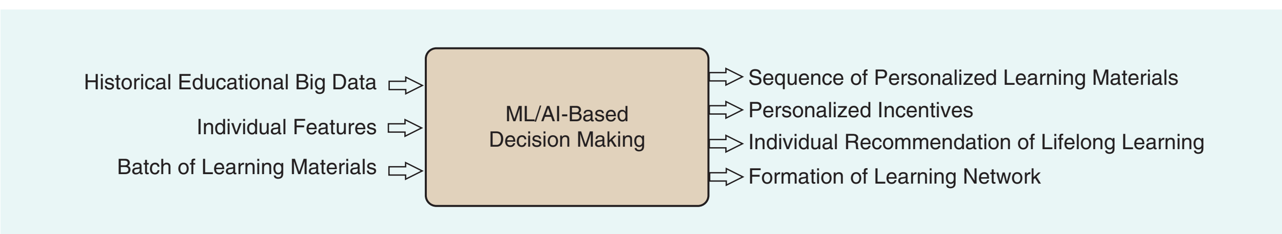

FIGURE 2. The basic concept of AI-empowered personalized education.

education remain unexplored, including both dark and bright sides. In this article, we discuss six core topics, review existing work, outline their limitations, and propose future research directions (see Figure 3 for an overview).

When discussing any form of education, *quality* is an inevitable keyword. The quality of education depends largely on the quality of the available learning content and on the quality of the personalized recommendations that guide each learner to the most suitable learning content. Thus far, researchers have studied the production of learning content, from developing AI-driven smart learning content such as intelligent, interactive textbooks and game-based learning platforms, to automatically generating learning content from the wild. The authors in [1], for example, develop a sentence-deletion method for text simplification. In [2], the authors investigate the effectiveness of discourse in multimedia to extract the knowledge from textbooks. Moreover, a large body of the existing work investigates the recommendation of both macro- and microlevel learning content, including courses in learners' degree plans as well as specific remedial content, such as lecture notes, videos, and practice problems. For example, in [3], the authors take advantage of a multiarmed bandit framework to optimize the selection of learning resources and questions to satisfy the needs of each individual student. Furthermore, another paper has developed an e-learning recommender system framework based on two concepts: peer learning and social learning, which encourage students to cooperate and learn jointly. Despite great effort, there remain several challenges to address, including content recommendation at heterogeneous levels, the recommendation of a bundle of related content followed by performance evaluation, and the Pareto-optimization of conflicting objectives in the content recommendation. We discuss these progress and future steps in the "Content Production and Recognition" section.

Historically, education has been tightly coupled with evaluation. In personalized education, assessment and evaluation concern both the learner's performance and the effectiveness of the intelligent learning platform. The early approaches that were

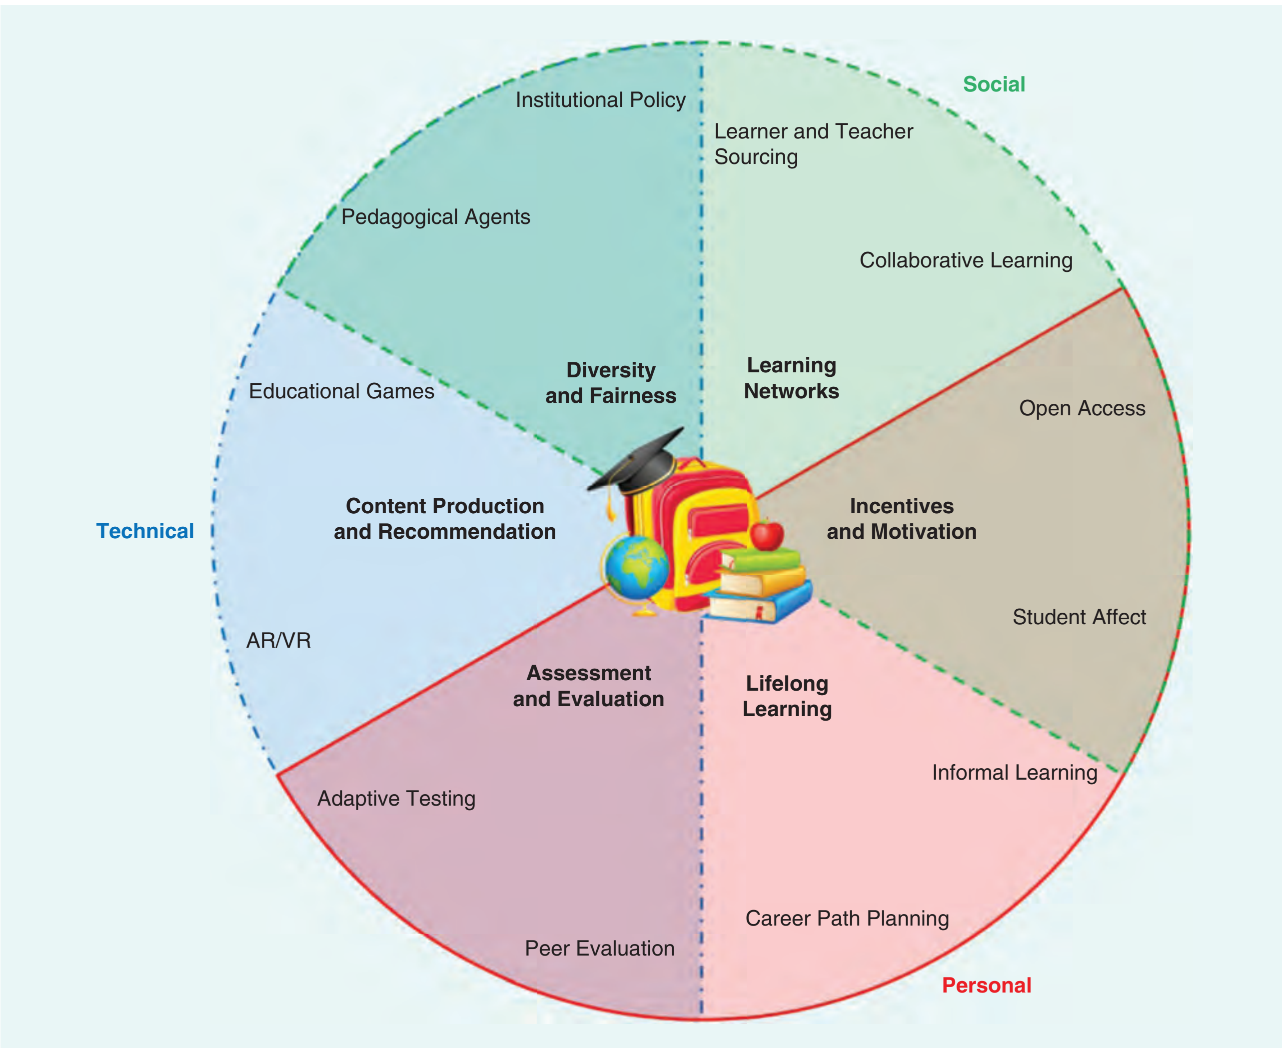

FIGURE 3. A list of (some of) the topics in personalized education, organized according to three different aspects: technical, personal, and social. In this article, we focus on six of the topics. AR/VR: augmented reality/virtual reality.

used for learner assessment, such as classical testing theories (CTTs), utilized graded standardized test summaries. Recent approaches include item response theory (IRT) models, which facilitate the estimation of latent-knowledge mastery levels, and knowledge tracing (KT) models, which pursue the evolution of a learner's knowledge. In [4], the authors compare CTTs and IRTs. Methods such as computerized adaptive testing improve the efficiency of assessments. The current approaches employed to evaluate learning platforms use rigorous experiments, often large-scale, randomized, and controlled trials. In this area, open problems include the prediction of learners' future performance, which enables providing better recommendations and more accurate feedback. This is referred to as the *KT* problem, for which several methods have been developed in the past few decades. As an example, among many others, some papers discuss a Bayesian framework for KT. Another challenge is to reduce the information loss while grading the arrived input from the learner through the accurate interpretation of knowledge level based on the test design. We elaborate on and address such challenges in the "Assessment and Evaluation" section.

Significant advances in science, technology, and health care have changed the working life of humans. Individuals have many more alternatives when choosing a job; he/she tends to change jobs more frequently than before, is more open to mobility, and has a long career span. As such, continuing education, which aims at advancing one's educational process, as well as lifelong learning, i.e., pursuing additional professional qualifications, are important components of educational policy in the world. Implementing these two concepts successfully has a significant impact on social welfare by developing new skills that enhance personal and professional lives.

During the past decade, AI-/ML-based personalized education has been under intensive investigation from several perspectives; nonetheless, the aforementioned aspects are largely neglected. Indeed, personalized education shall accompany the learner throughout his/her life, which can be difficult and costly to implement. Other challenges include a lack of appropriate data, a potentially long delay to receive feedback, high diversity, and fast dynamics in the environment. For instance, in [5], the authors design a new genre of educational technology, personal computer systems, which support learning from any location throughout a lifetime. Another research direction is to enable the learning system to learn continuously. Parisi et al. [6], for example, investigate the ability of neural networks to provide lifelong learning. We elaborate more on this topic in the "Lifelong Learning" section.

Similar to any other task, humans require motivation for learning. Generally, incentives for learning can be defined as *an inducement or supplemental reward that serves as a motivational device for intended learning* [7]. Presumably, the most conventional models of incentive are the grade and the certificate, which are implemented as a part of learning platforms to motivate students. The strength of such motivation depends on the validity and acceptance of such certificates by different authorities, such as employers. Nonetheless, employing AI methods for personalized education enables incentive design far beyond the issuance of a certificate. This includes monetary rewards in the form of bonuses for online learning materials. The incentives can also be introduced using soft methods, such as gamification based on the learner's characters to promote continuous learning, or by adapting the features of the learning environment based on the learner's traits to engage him/herself in the learning process for as long as possible. In [8], the authors investigate the effects of gamification on students' motivation from several perspectives. We discuss these challenges and methods in the "Incentives and Motivation" section.

Education is social, and learners can benefit meaningfully from their peers. It is therefore urgent to develop effective ways to build networks that serve as a conduit of knowledge for learners to interact with each other. In its current form, personalized education suffers from a lack of student–student and student–teacher connections and interactions, which, unquestionably, have a positive impact on learning, through discussions, joint efforts, and brainstorming. For example, Vesely, et al. [9] investigated and compared the influence of such communities from both students' and teachers' perspectives. As another example, in state-of-the-art research, authors have studied the building and sustaining community in asynchronous learning networks, i.e., when the learners are physically separated.

Despite past research efforts, we believe that by capitalizing on AI and ML methods, online platforms have more to offer, especially for building the knowledge and expertise networks that facilitate the assimilation and dissemination of information, and, consequently, by enabling close interactions (in terms of mentorships, friendships, coworkers, and so on), creating knowledge. Personalized education platforms can promote autonomous network formation by encouraging learners to interact. Moreover, the platforms can establish links among those learners that satisfy some similarity conditions and can hence be useful to each other for cooperation, inspiration, and motivation. Still, it is vital to note that online contacts can be lost easily, and the learners, especially at early ages, are more prone to feeling isolated and depressed. We elaborate on these issues in the "Building Learning Networks" section.

In many different ways, education affects the well-being of humans, and society at-large, both in the short and long term. As such, fairness is a highly important aspect of education, regardless of whether it occurs in conventional classrooms or in modern platforms that can personalize the learning experience. Despite this great importance, personalized education, similar to its traditional counterpart, might result in and strengthen inequality. This arises, for example, due to unequal access to learning platforms, biases in training data, inaccuracies in algorithm design, and so forth. Indeed, existing research shows that some subgroups of students—many of whom are also privileged with respect to conventional education platforms—would profit more from personalized education than would their peers. To address this issue more specifically, there have been intensive efforts to develop appropriate fairness models [10]. Moreover, several research works have studied the fairness of predictive algorithms in educational settings.

Another crucial issue is diversity. Today, it is well established that diversity promotes innovation and efficiency in the working place. Nonetheless, given the social responsibility of education, recruiting only diverse talents does not suffice. An AI-based personalized education platform can help diversify the education environment, for example, by rewarding collaborative learning in diverse networks. We discuss these topics in the "Diversity, Fairness, and Biases" section.

Content production and recommendation

Ultimately, the quality of education depends on the quality of the learning content. Creating new content requires the wisdom of human content designers and educational experts; to date, AI methods have not shown the capability of creating learning content on their own. However, they still have plenty to offer in content production by automating mundane jobs and helping humans in tasks where human input is necessary. Specifically, the role of AI should be to 1) take away repetitive tasks that can be automated and 2) assist humans by providing the feedback extracted from data during the process of content production in a human-in-the-loop manner. The following section detail ample future research directions in content production.

Content summarization and question generation

In many educational domains, knowledge is factual. For example, in history, one often needs to remember specific detailed facts about historical events. Even in scientific domains such as biology, there is also factual knowledge, e.g., the size and life span of an animal. In this case, there are many natural language processing (NLP)-based tools that can be used for content production. For example, text summarization tools can sort through long, sometimes redundant, textbook sections and extract key facts for remedial studies. This is not only helpful but also sometimes crucial to certain learner groups, such as those with learning disabilities. Moreover, automatic question generation can effectively produce high-quality factual assessment questions that have short, textual answers [11]. An example of automated question generation is shown in Table 1: We reversed a long short-term memory network-driven question-answering pipeline trained on common question-answering data sets, turned it into a question-generating pipeline, and applied it to textbooks. Human experts have indicated that the quality of generated questions is higher than those produced from other methods [11].

Multimodal content understanding

Many educational domains involve multimodal learning content, such as text, formulas, figures, and diagrams. When a learner fails to answer an assessment question correctly, personalized education systems need to automatically retrieve relevant content to help the learner resolve his/her confusion (by retrieving examples and explanations) or give the learner more practice opportunities (by retrieving assessment questions). Retrieving content within the same modal is relatively easy; for example, when a learner answers a textual question incorrectly, it is possible to use information-retrieval methods to extract relevant textbook chunks or lecture slides. However, when the most helpful content is in another modality, such as when a Venn diagram is the most effective at helping a learner clear up a misconception in a probability question involving text and mathematical formulas, it is hard to retrieve the diagram. Therefore, more work needs to be done when the domain includes multimodal content. To understand these content modalities and use them for content production, we need to learn universal representations across all modalities, possibly using embedding approaches to map multiple modalities into a shared vector space.

Human-in-the-loop content design

Even for humans, learning content is not created in one shot. Similar to textbooks, which have different editions, learning content is frequently edited and updated over time. Therefore, during this multistep process, we can use AI methods to act as (possibly even interactive) assistants to content designers. The duties that can be assigned to AI methods include 1) analyzing learners' data to identify the areas of priority for new content and assessment questions that need to be improved (see the "Assessment and Evaluation" section for discussions on how existing learner assessment models can also provide information about question quality), 2) providing drafts of instructor responses and perform automated checking of human-generated content using NLP tools, and 3) using crowdsourcing to put the learning content together by soliciting on-demand feedback. The third task is especially important in online educational settings, where learning occurs during frequent exchanges between learners and human instructors and assistants [12].

Even with high-quality learning content, the presentation, i.e, the personalized recommendation of the right learning content to the right learner at the right time is crucial to optimizing learning outcomes. Fortunately, this is an area at which AI methods can excel: By automatically deploying recommendations and analyzing the data of learners' performance, they can quantify the effect of the learning content on certain learners in terms of specific learning outcomes to detect the most effective ones. On the contrary, humans, including educational experts in the past, use theoretical models of learning and do not fully take advantage of this data. In the following sections, we discuss a few directions for future research in this area.

Recommendations at the microscopic and macroscopic levels

Learning content is organized at multiple levels, down to individual paragraphs and assessment questions, up to courses and

Table 1. An example of two automatically generated assessment questions for two different answers with the same input context from a textbook. The answers are underlined and marked with different colors in the input context.

Context (Biology): On each chromosome, there are thousands of genes that are responsible for determining the genotype and phenotype of the individual. A gene is defined as *a sequence of DNA that codes for a functional product.* The human haploid genome contains **3 billion base pairs** and has **between 20,000 and 25,000** functional genes.

3 billion base pairs	Between 20,000 and 25,000
How many base pairs are on the human genome?	How many functional genes are on the human haploid genome?

textbooks that organize several pieces of learning content together. Therefore, we need to study content recommendations at multiple levels: 1) microscopic, which comprises individual questions and lecture video suggestions [13], and 2) macroscopic, which includes course recommendation, especially for learners taking MOOCs [14].

Efficient experimentation and synthetic learner models

Traditionally, the fields of learning science and education have relied on rigorous A/B testing to validate the educational impact of learning content, usually in terms of its ability to improve learning outcomes for learners in the experimental group over those in the control group. However, this approach leads to long experimental cycles because 1) typically, only one learning content at a time can be validated and 2) metrics such as long-term learning outcomes naturally require long experimental cycles. Therefore, it is imperative to search for novel tools that enable rapid experimentation. Possible solutions include employing Bayesian optimization to test multiple contents simultaneously, or utilizing reinforcement learning (RL) as more and more learners use a piece of learning content. In the past, using RL to learn instructional policies (content recommendation can be viewed as a form of the instructional policy) has been limited due to a lack of large-scale real learner data; however, recent approaches have looked at using data- or cognitive theory-driven synthetic learner models to simulate learner data.

Conflicting objectives

There is no unified objective in personalized learning because learning outcomes themselves are defined at multiple time scales. The optimal action may differ across different objectives. For example, the learning content used in a practice session that maximizes a learner's performance on the midterm exam tomorrow may differ from the one that maximizes their overall course grade, which may differ from the one that maximizes their chance of getting a specific job after graduation. Therefore, we need to develop personalization algorithms that can balance multiple objectives and even resolve potential conflicts among these objectives. We also need to understand how these objectives interact with each other; for example, what skills taught in courses and schools carry over after graduation—a key issue in lifelong learning (discussed in detail in the "Lifelong Learning" section).

Figure 4 depicts the interplay between different elements, such as context, prediction, feedback, and so on to optimize course recommendation. It is worth noting that the approaches described previously are generic in the sense that they have wide applicability to different educational areas, including signal processing, possibly with minor domain-dependent adaptations. As an example, in [15], the authors apply several of the aforementioned ideas to develop eTutor, a personalized, web-based education system that learns the optimal sequence of teaching materials to present, based on the student's context

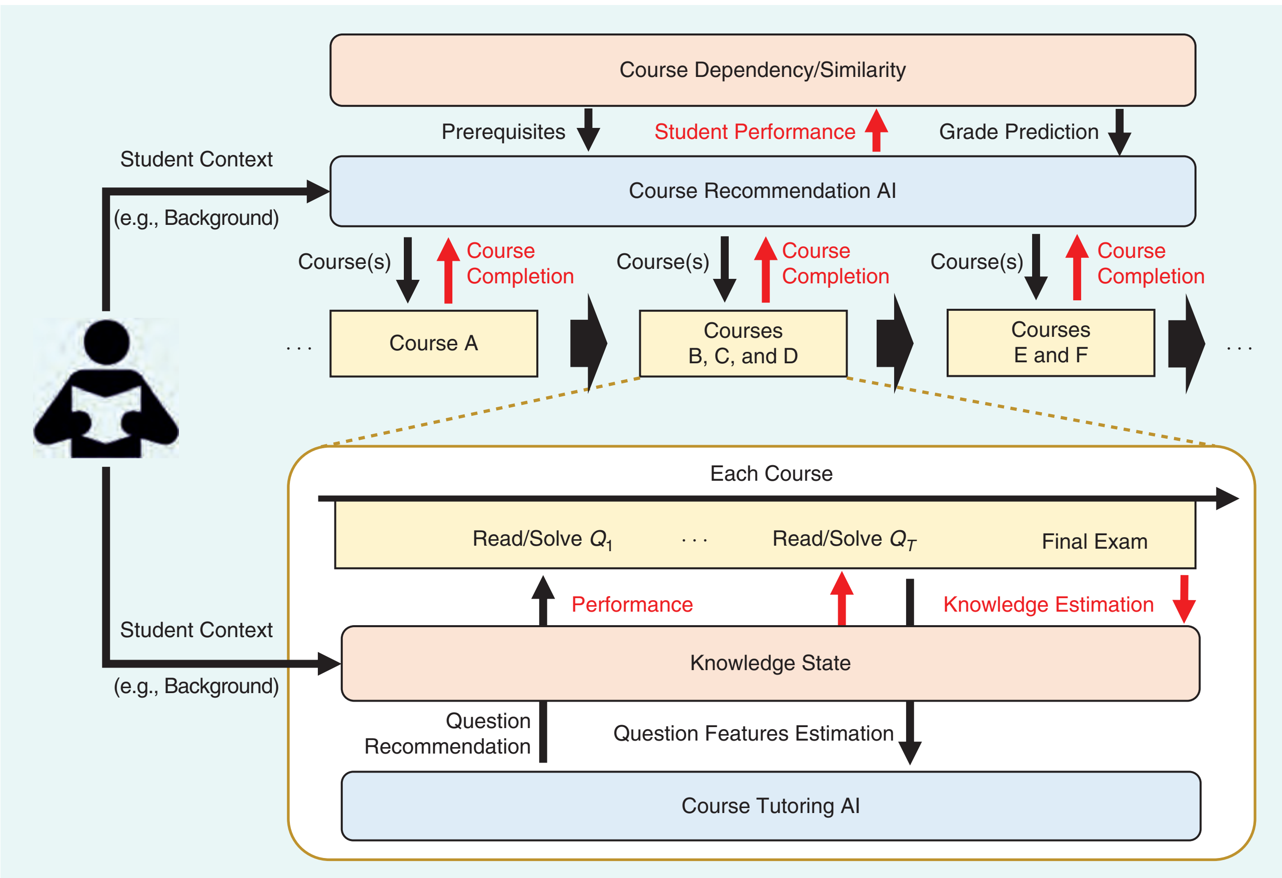

FIGURE 4. A detailed framework of course recommendation.

and feedback, previously shown teaching materials. In an experiment, they apply the eTutor system in the following scenario: The students have studied digital signal processing in the past. The role of eTutor is to recommend learning materials to the students with the goal of refreshing their minds about discrete Fourier transform in the least amount of time. The eTutor shows better performance compared to random- and fixed-selection rules.

Assessment and evaluation

A key problem in learner assessment is estimating how well he/she will master each knowledge component/concept/skill from the responses to assessment questions. Related works can be broadly classified into two model categories: 1) static, which analyze the generated data as learners take an assessment (assuming that each learner's knowledge remains constant during the assessment) and 2) dynamic, which track a learner's progress throughout a (possibly long) period as their knowledge levels evolve. The following is a short overview of each category:

- *Static models—IRT*: The basic 1PL IRT model characterizes the probability that a learner answers a question correctly as

  $$P(y_{i,j} = 1) = \sigma(a_j - d_i),$$

  where $y_{i,j}$ denotes the binary-valued graded response of learner j to question i, where 1 implies a correct response and 0 otherwise. Moreover, $a_j \in \mathbb{R}$ and $d_i \in \mathbb{R}$ are scalars that correspond to the learner's ability and the question's difficulty, respectively. Also, $\sigma(\cdot)$ is a link function that is usually the sigmoid function or the inverse probit link function. Later extensions include 2PL IRT models, which add another multiplicative scaling parameter. This parameter corresponds to the ability of each question, differentiating high-capacity learners from low-capacity ones. Further, 3PL IRT models add another scalar outside of the link function, which corresponds to the probability that an item can be guessed correctly. Finally, multidimensional IRT models use vectors instead of scalars to parameterize strengths and weaknesses to capture multiple aspects of one's ability. Using the aforementioned models, one can 1) obtain relatively stable estimates of learners' ability levels by denoising learners' responses and 2) estimate the quality of each assessment question.
- *Dynamic models—KT*: KT models consist of the learner performance $[f(\cdot)]$ and learner knowledge evolution models $[g(\cdot)]$, expressed as

  $$y_t \sim f(a_t), \quad a_t \sim g(a_{t-1}),$$

  where t denotes a discrete-time index. Early KT models, such as Bayesian KT [16], treat knowledge (h_t) as a binary-valued scalar that characterizes whether or not a learner masters the (single) concept covered by a question. Performance and knowledge evolution models are simply noisy binary channels. Later, factor-analysis-based KT models used a set of handcrafted features, such as the number of previous attempts; successes; and failures on each concept, to represent a learner's knowledge levels. These models require expert labels to associate questions to concepts, resulting in excellent interpretability because they can effectively estimate the knowledge level of each learner on expert-defined concepts. Recent KT models incorporate deep learning, especially recurrent neural networks into the KT framework, where knowledge is represented as a latent vector $\mathbf{a}_t$. These models achieve state-of-the-art performance in predicting future learner responses, although in some cases, the advantage is not significant despite the loss of some interpretability.

The existing learner assessment models have several bottlenecks. First, there are not many models with both the ability to achieve state-of-the-art performance in data fitting, (i.e., future performance prediction) as well as feedback generation (i.e., providing interpretable feedback to learners and instructors for downstream tasks, such as personalization). Therefore, it is imperative to develop new deep learning-based models that not only inherit the flexibility of neural networks to accurately predict learner performance but also build in cognitive theory-inspired structures to promote interpretability and enable the generation of meaningful feedback. As an example, in the recently developed attentive KT (AKT) model [17] (Figure 5), we combined state-of-the-art attention networks with cognitive theory-inspired modules. We used a monotonic attention mechanism where weights exponentially decay over time and question embeddings are parameterized by the 1PL IRT model to prevent overfitting. Experimental results show that the AKT model not only outperforms existing KTs but also exhibits some interpretability [17]. The existing learner assessment models operate almost exclusively on graded learner responses; however, converting raw learner responses to graded responses leads to considerable information loss.

For multiple-choice questions, different distractor options are not created equal; choosing certain incorrect options over others might indicate that a learner exhibits a certain misconception. However, this information is lost when the learner's option choice is converted to a graded response. Moreover, due to their superior pedagogical value, open-response questions have been widely adopted; the specific open-ended response a learner enters contains rich information about his/her knowledge state. Therefore, it is vital to develop models that go deeper than the graded response level and into the raw response level. These models allow for personalization at even finer levels; for example, after each step as a learner solves an open-ended mathematical problem step by step and by enabling personalized education systems to attend to learner difficulties in a more timely manner.

Another consideration in effective learner evaluation is that assessment and performance prediction models must be tailored to different learning environments and platforms. For instance, the accurate prediction of students' future college performance based on their ongoing academic records is crucial to carrying out effective pedagogical interventions so that on-time, satisfactory graduation is ensured. However, foretelling student performance in completing degrees (e.g., college programs)

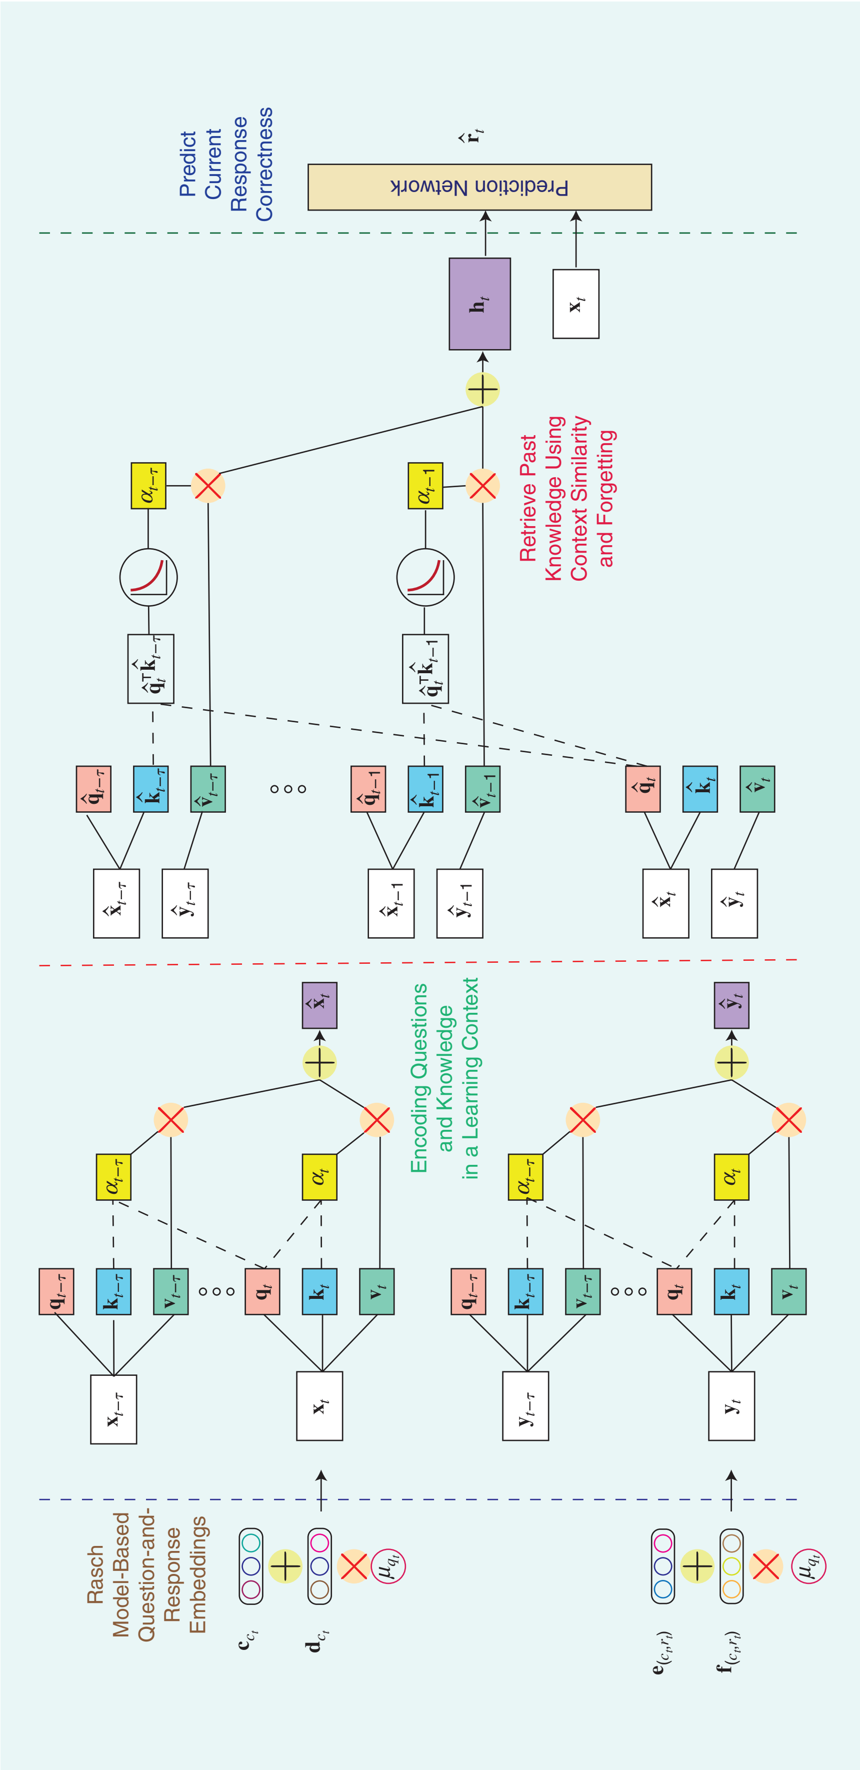

FIGURE 5. An overview of the AKT method. We used IRT-based raw embeddings for questions and responses. We computed context-aware representations of questions and responses using two encoders. We then used a knowledge retriever to retrieve past-acquired knowledge for each learner using a monotonic attention mechanism, which is used for performance prediction.

is significantly different from that for in-course assessment and intelligent tutoring systems. The following describes the most important reasons for why this is so.

- First, students differ tremendously in terms of backgrounds as well as the study domains (majors, specializations), resulting in different selected courses. Even if the courses are similar, the sequences in which the students take the courses might differ significantly. Therefore, a key challenge for training an effective predictor is to handle heterogeneous student data. In contrast, solving problems in intelligent tutoring systems often follow routine steps that are identical for all students. Similarly, predictions of students' performance in courses are often based on in-course assessments that are identical for all students.
- Second, although the students often take several courses, not all of them are equally informative for predicting the students' future performance. Utilizing the student's past performance in all courses that he/she has completed not only increases complexity but also introduces noise in the prediction, thereby degrading the performance. For instance, while it is meaningful to consider a student's grade in linear algebra for predicting his/her grade in linear optimization, the student's grade in chemistry lab may have much weaker predictive power. However, the course correlation is not always as obvious as in this example; therefore, to enhance the accuracy of performance predictions, it is essential to discover the underlying correlation among courses.

- Third, predicting student performance in a degree program is not a one-time task; rather, it requires continuous tracking and updating as the student finishes new courses over time. An important consideration is the following: The prediction shall be made based on not only the most recent snapshot of the student's accomplishments but also the evolution of the student's progress, which may contain valuable information to improve the prediction's accuracy. However, the complexity can easily explode because just mathematically representing the evolution of student progress itself can be a daunting task. Treating the past progress as equally as the current performance when predicting the future may not be a wise choice either because old information tends to be outdated.

Finally, we would like to emphasize the following: Similar to an offline system, in AI-powered personalized education, an assessment does not remain limited to evaluating the performance of individual students in different tests in a single online education portal. Indeed, evaluation might be necessary not only for individuals but also for a collection of students as well as other stakeholders, such as educators, policy makers, and the providers of online education. In particular, the fair and precise comparison, analysis, and accreditation of online education portals, as well as the degrees and certificates provided by such portals, are crucial. The reasons include the following: 1) Distance education has grown into a broad industry in the past decade, 2) the majority of learners rely on certificates of online classes as approval for obtaining the necessary knowledge and skills, and 3) online education is inherently international and crosses boundaries. Similar to improving the evaluation of students, AI and ML methods, together with bid data analysis, can assist in the accreditation and comparison of online portals and the degrees and certificates issued; this includes, e.g., comparing the average student's performance with an online degree to that of a traditional, yet accredited, degree. A detailed discussion of such topics has several perspectives and is therefore beyond the scope of this article.

Lifelong learning

Lifelong learning emphasizes holistic education and the fact that learning takes place on an ongoing basis as a result our daily interactions with others and with the world around us in different contexts. These include not only schools but also homes and workplaces, among several others. Because of its ongoing nature, making foresighted learning plans is crucial for lifelong learning to achieve the desired outcome.

In the school context, a specific challenge for developing a learning plan is the course sequence recommendation in degree programs [14]. Recent studies show that the vast majority of college students in the United States do not complete college in the standard time frame. Moreover, today, compared to a decade ago, fewer college students graduate in a timely manner. Although several factors contribute to students taking longer to graduate, such as credit losses resulting from a school transfer, uninformed choices due to low advisor-student ratios, and poor preparation for college, the inability of students to attend the required courses is among the leading causes. If a student selects courses myopically without a clear plan, he/she may end up in a situation where required subsequent courses are offered (much) later, thereby (significantly) prolonging the graduation time. Hence, to accelerate graduation, students should essentially select courses in a foresighted way while taking the course sequences (shaped by courses being mandatory, elective, or prerequisite) into account.

It is also essential to observe the time period in which the school offers various courses. More importantly, as the number and variety (in terms of backgrounds, knowledge, and goals) of students is expanding rapidly, the same learning path is unlikely to best serve all students. Therefore, it is crucially important to tailor the course sequences to students. To this end, it is necessary to learn from the performance of previous students in various courses/sequences to adaptively recommend course sequences for current students. Obviously, this depends on the student's background and his/her completion status of the program to maximize any of a variety of objectives, including the time to graduate, grades, and the tradeoff between the two. To make such plans, AI is a tool of great potential; however, designing AI technologies for personalized, foresighted, and adaptive course planning is challenging in several dimensions, as described in the following.

- First, course sequence recommendation requires dealing with a large decision space that grows combinatorially with the number of courses.
- Second, there is a great deal of flexibility in course sequence recommendation as multiple courses can be taken simultaneously, while it is also subject to many constraints due to prerequisites and availability.
- Third, any static course sequence is suboptimal as the knowledge, experience, and performance of a student develops and evolves during the process of learning.
- Finally, students vary tremendously in their backgrounds, knowledge, and goals.

For example, in [14], we develop an automated course sequence recommendation system to address the aforementioned challenges. To reduce complexity and enable tractable solutions, we solve the problem in two steps, as illustrated in Figure 6:

1) The first step corresponds to offline learning, in which a set of candidate recommendation policies are determined that minimize the expected time to graduation or the on-time graduation probability using an existing data set of anonymized student records based on dynamic programming.
2) The second step corresponds to online learning, in which for each new student, a suitable course sequence recommendation policy is selected depending on this student's background using the learned knowledge from the previous students.

In other lifelong learning contexts (e.g., the workplace), although similar challenges may still be present, new challenges are likely to emerge, and hence, foresighted learning plans must be tailored to the specific context.

Recent research shows a significant gap between the lectures offered in schools and job requirements, especially in emerging disciplines like data science. Soft skills such as communication and teamwork are often even more important than typical technical skills. Future research on lifelong learning shall bridge this gap. Indeed, there is a systematic demand for the research community to identify and study the skills that significantly contribute to professional perspectives instead of maximizing achievement in schools. Educators can take advantage of the findings to adjust school curricula and educational activities to better prepare students for the future. The centerpiece of possible approaches is to fuse a student's school records with future employment outcomes, possibly tracked over a long period, as well as other data sources such as course syllabi and job postings, to identify the crucial skills that extend from the classroom (virtual or otherwise) to the profession. There is also a need for research labor studies to conduct interviews with 1) employers, to understand their requirements; 2) job seekers, to identify the skills to acquire; and 3) training providers, to clarify the skills that can be taught in a part-time or on-the-job way rather than through centralized educational programs, given workers' real-life constraints.

Incentives and motivation

Thus far, one crucial aspect of personalized education has been largely left aside, namely, motivation and incentive design. This is unfortunate as these factors significantly contribute to a learner's perseverance and engagement, and thereby, overall student achievement. As such, they affect not only individuals but also the entire society in terms of the efficiency of resource expenditure.

In educational sciences, motivation is regarded as a concept that involves several learning-related features, such as initiation, goal orientation, intensity, persistence, and the quality of behavior [7]. Therefore, as described by Hartnett in [18], *motivation is an unobservable dynamic process that is difficult to measure directly, but it is inferable from observations*. Similar to the other crucial factors of successful education, such as talent and interest, motivation originates and is influenced by personal factors, including goals and beliefs. As such, it is reasonable to conclude that intelligent personalization affects motivation to a large extent.

Motivation can be intrinsic or triggered by external factors. Accordingly, the various features of personalized education, such as recommending a proper series of content or creating educational networks, might implicitly improve a learner's motivation by increasing the engagement level. Such efforts make the learning experience more pleasant, thereby improving a learner's satisfaction level. This is, however, insufficient. It is imperative to integrate direct motivating methods into personalized education and the learning platform. To this end, in the following sections, we describe a few frameworks that can accommodate motivation and its relevant concepts appropriately (see [18] for more information).

Behavioral economics

Any personalized education platform should be able to appropriately connect, interact, and interface with humans. Hence, proper operation significantly depends on various characteristics of the members of the target group that shape their decision making behavior. Indeed, a utility function is the most seminal computational model for the interests of learners. For a rational decision maker, the utility function is conventionally increasingly concave and is to be maximized. However, humans often demonstrate unusual patterns in their utility functions and decision making due to the following reasons:

- Humans make mistakes, often due to inaccurate beliefs and imprecise predictions.
- Humans often act irrationally and based on heuristics.
- Humans think and act in different ways as a result of their unique backgrounds, including personality and experiences [19].

Behavioral economics accommodates and formalizes such aspects; therefore, one can take advantage of behavioral economics for efficient incentive design and motivation in learning platforms [20].

Self-determination theory

This theory asserts that humans have an intrinsic urge to be self-autonomous, competent, and connected with respect to their environments [21]. Although behavioral economics is appropriate for investigating motivations that result from external rewards, self-determination observes motivations from an internal perspective. Indeed, any environment, including the learning platforms that satisfy the aforementioned needs of humans, awakens the intrinsic motivation, rendering external triggers mostly unnecessary. As such, promoting intrinsic

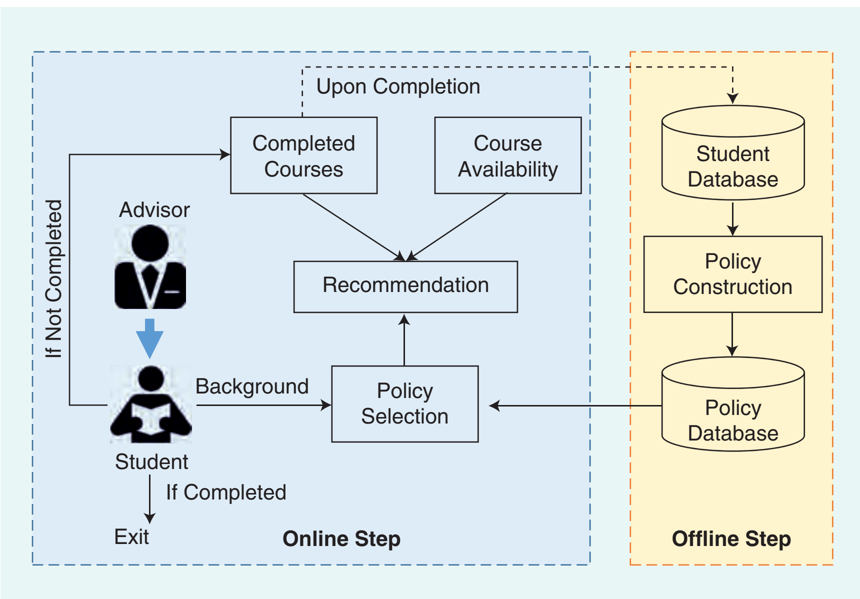

FIGURE 6. The course sequence recommendation.

motivation is significantly more effective than extrinsic motivation as it is often associated with lower cost when compared to material rewarding and has a long-lasting effect [18].

Self-efficacy theory

This concept corresponds to an individual's confidence in his/her capability of performing a specific task to be undertaken; for example, learning in an online learning platform or performing at a certain level [22]. Researchers show that humans constantly assess their self-efficacy, mainly based on the observed information from the environment and past experiences [22]. Similar to the self-determination theory, self-efficacy considers the intrinsic motivation, implying that a feeling of efficacious triggers the internal motivation feeling in learners. Other relevant concepts include interest and goal orientation [18].

The main challenge is how best to utilize AI and ML to motivate the learners of a personalized learning platform, based on the aforementioned theories that formalize and explain human behavior. To clarify this, consider the utility function of a learner in a personalized learning platform as an example [20]. The function quantifies the learner's well-being while using the platform, and, consequently, his/her (future) engagement. Some learners exhibit hyperbolic preferences, overweighting the presents ones so much so that future rewards are largely ignored. Some learners show strong reactions even to nonmonetary rewards, while other learners demonstrate reference-dependent preferences, implying that the utility is largely determined by its distance from a reference point; for example, a predefined goal or the average performance. By using ML and AI methods, a learning platform can take advantage of the available data and a learner's feedback to estimate the utility function of that learner, thereby predicting his/her reaction to the potential triggers of incentive and motivation. Consequently, the platform can adjust and allocate the reward among learners efficiently and fairly.

As another example, consider the self-determination theory. Based on this theory, a sophisticated personalized learning platform guarantees choice, connectedness, and the feeling of competence for the learner. To this end, the design of recommendation tools based on AI and ML methods should allow for enough alternatives, both at the micro- and macrolevels, to ensure autonomy. Moreover, the suggested learning content should be based on the learner's feedback and the results of accurate assessment to avoid inducing a feeling of incompetence in the learner. In addition, promoting network formation or establishing a link between coherent learners and intensive interaction results in connectedness. This is also in accordance with the self-efficacy theory, in the sense that by providing appropriate feedback and suitable side information, the platform increases the positive belief of a learner in his/her ability to perform well on a learning platform.

Building learning networks

A potential negative effect of personalized education, especially in an online environment, is a loss of peer interactions and of the sense of community that are usually present in traditional classrooms. Fortunately, the rise of online social networks seems to facilitate interaction and networking between teachers and learners, as does the coproduction of content both inside and outside the classroom. Learning applications and pedagogy can also be built based on online social networks to bridge formal and informal learning as well as to promote peer interactions on both curricular and extracurricular topics. Moreover, various education-related social networks have been created to facilitate collaboration, post/answer questions, and share resources; however, a formal method to build these learning networks and a deep understanding of their effectiveness are absent.

The core of learning networks is peer interaction, which has important implications for personalized education when teaching resources are limited. For example, peer review serves as an effective and scalable method for assessment and evaluation when the number of students enrolled in a course far exceeds the number of teaching assistants; however, in learning networks, effective peer review poses new challenges [23]. On the one hand, peer reviewers have different intrinsic capabilities, which are often unknown. On the other hand, peer reviewers can choose to exert different levels of effort (e.g., time and energy spent in reviewing), which is unobservable. Identifying unknown intrinsic capabilities corresponds to the adverse selection problem in game theory. A natural candidate for solving this problem is to use matching mechanisms, i.e., assign reviewers to students. Existing works on matching mechanisms focus on one-shot peer interactions and design one-shot matching rules. However, their assumptions do not hold in peer-review systems, where the review quality depends crucially on the reviewers' effort.

Motivating reviewers to exert high effort corresponds to the moral hazard problem in game theory. One way to address this problem is to use social norms, where each peer reviewer is assigned a rating that summarizes his/her past behavior and recommends a "norm" that rewards a reviewer with good ratings and punishes those with bad ratings. However, existing works on social norms assume that peer reviewers are homogeneous. This assumption does not hold in peer-review systems because different reviewers have different intrinsic capabilities. Because a peer reviewer's ultimate review quality is determined by his/her intrinsic capabilities and effort, designing effective peer-review systems in learning networks becomes significantly more challenging due to the presence of both adverse selection and moral hazards. Therefore, new peer-review system designs should simultaneously solve both problems so that peer reviewers find it in their self-interest to exert high effort and receive ratings that truly reflect their capabilities.

Another primary function of learning networks is to foster learning content coproduction and sharing. Building such learning networks is vastly different from building traditional networks (e.g., computer and transportation networks); as with learning networks, individual learners create and maintain the links. Because links permit the acquisition and dissemination of learning content, it is theoretically intriguing and practically valuable to have a deeper understanding of the networks that are more likely to be formed by self-interested learners. Game theory is a useful tool to formulate and understand the strategic behavior of learners. The formulation must capture the heterogeneity of learners in terms of goals, capabilities, costs, and self-interest nature [24]; that is, each learner intends to maximize his/her benefit from content coproduction and sharing, minus whatever the cost is to establish the links.

Our previous work [25] studies the endogenous formation of networks by strategic, self-interested agents who benefit from producing and disseminating information. The results showed that the typical network structure that emerges in equilibrium displays a core-periphery structure, with a smaller number of agents at the core of the network and a larger number of agents at the periphery of the network. Furthermore, we established that the typical networks that emerge are minimally connected and have short network diameters, which are independent of the size of the network. In other words, the theoretical results show that small diameters tend to make information dissemination efficient, and minimal connectivity leads to minimizing the total cost of constructing the network. These results are consistent with the outcome of numerous empirical investigations. Such theoretical analyses and tools are essential guides for building learning networks. Also, based on this analysis, one can create protocols to motivate selfish learners to take actions that promote systemwide utility.

Future research into learning networks hinges on understanding the knowledge flow between students via peer interaction. Such an understanding enables educators to effectively moderate peer interactions and to encourage the interactions that promote peer learning. Peer learning is especially valuable as education extends into more diverse settings, such as remote online learning during the COVID-19 pandemic. In these settings, it is difficult for instructors to moderate learning activities remotely; hence, peer learning through online course discussion forums becomes essential. It is therefore vital to understand

- interaction tendencies and students' behavior in these discussion forums [26]
- the flow of knowledge by combining discussion forum activities with grades
- the factors that enhance knowledge flow
- the design of automated strategies that moderate student activities when necessary.

Diversity, fairness, and biases

Experimental studies show that AI-driven personalization such as student assessment, feedback, and content recommendation improve overall learning outcomes; nonetheless, certain student subgroups may benefit more than other subgroups due to the biases that exist in training data [27]. This imbalance jeopardizes students who are already underserved, particularly because they often have limited access to advanced, digitized educational systems and are infrequently represented in the data sets collected by these systems [28]. As a result, it is essential to develop AI tools that promote fairness among learners with different backgrounds, thereby making education more inclusive for future generations.

To mitigate biases and to promote fairness and equity in AI methods, currently, researchers pay significant attention to developing approaches that promote fairness, primarily in the context of predictive algorithms:

- The first major research problem studied is how to properly define fairness. Many definitions of fairness exist, including individual fairness, which requires that users with comparable feature values be treated similarly; parity in the predicted probability of each outcome across user groups (drawn using sensitive attributes); parity in the predicted probability of each outcome given actual outcomes regardless of sensitive attributes; and counterfactual fairness, which requires that the predicted outcome for each user remains mostly unchanged if the sensitive attribute changes.
- The second major research problem concerns developing methods that enforce fairness in predictive algorithms. Existing approaches include preprocessing the data to select only the fair features as inputs to algorithms, and postprocessing the output of algorithms to balance across user groups. The most promising approach is to impose regularizers and constraints while training predictive algorithms. These methods result in better fairness at the expense of sacrificing some classification accuracy; however, they are empirically shown to obtain better tradeoffs between fairness and accuracy than other fairness-promoting methods.

Promoting fairness and equity is a necessity of education that requires a comprehensive approach for it to be fulfilled: We need to not only design fair personalization algorithms but also develop systematic principles and guidelines for their application in practice. In other words, we need a set of tools to regulate the use of AI algorithms.

Finally, despite its great promise, AI-driven personalization in education can also bring risks that have to be closely monitored and controlled. Recently, there have been calls for a U.S. Food and Drug Administration-type framework for other AI applications, such as facial recognition. It is essential to establish a similar ecosystem in education with a set of regulations around the issues of data ownership, sharing, continuous performance monitoring, and validation to control every step of the process, from ensuring the diversity and quality of the collected data, developing algorithms with performance guarantees across different educational settings, to identifying misuse and implementing a fail-safe mechanism.

COVID-19 and AI-enabled personalized education

Among its several other adverse effects, the COVID-19 pandemic has disrupted or interrupted the functionality of conventional education systems around the globe. Not surprisingly, students have experienced this adverse effect to varying degrees depending on several factors, including country/region, family status, and individual characteristics. The complications differ over a wide spectrum and include reduced learning ability, depression, loss of concentration, and a decline in physical fitness. Such issues arise mainly due to spending less or no time at school, where students receive educational materials and support in learning, interact with their peers and teachers, develop incentives, and are evaluated. Furthermore, many students cannot take full advantage of replacement resources such as online materials, e.g., in the absence of an appropriate technological device/a reliable Internet connection or a suitable learning environment at home. As a consequence of its vitality, the impact of COVID-19 on education has attracted a great deal of attention. For example, in [29], the authors describe the influence of pandemic-triggered growth in online learning on students' performance and equity. Several other works, such as [30], study the domain-specific educational effects of the pandemic, evaluate available solutions, and provide suggestions for policy makers to compensate for the pandemic's negative educational consequences.

Personalized and distance education had already been trending upward in the past decade. Still, the COVID-19 pandemic has urged both public and private sectors to rapidly increase investigations into R&D in this area to earn individual and/or social profit. For instance, the pandemic has increased the use of online learning tools for signal processing education, especially at the undergraduate level. These include web-based laboratories for digital signal processing [31] and online ML education modules [32]. Although it is essential to carefully study this tremendous push toward revolutionizing education from several perspectives, in the scope of our article, we confine our attention to the role and influence of AI and ML.

As described previously, AI and ML have great potential to enhance online education in different ways, e.g., through improving the quality of learning materials, enabling fairness and diversity, generating proper tests, and allowing for the construction of knowledge networks. The latter way is a universal aspect of applying AI and ML methods in distance and asynchronous education regardless of the current pandemic; nevertheless, such methods can additionally assist in accelerating the rebuilding of education systems and in mitigating the pandemic's detrimental effects. For example, by using ML methods on the available data, policy makers can classify students based on their exposure to the educational effects of a pandemic; using this classification, one can allocate resources efficiently while satisfying fairness constraints. As another example, by taking advantage of ML methods, one can optimize a school's closure plan based on different features, such as neighborhood, size, grade, and so forth.

Summary and conclusions

Enabling personalized education is one of the most precious merits of AI, relative to education. This paradigm significantly improves the quality of education in several dimensions by adapting to the distinct characteristics and expectations of each learner, such as personality, talent, objectives, and background. Besides, online education is of the utmost value under abnormal circumstances, such as the COVID-19 outbreak or natural disasters. Indeed, conventional education requires significantly more resources than the online format with regard to education space, scheduling, and human resources, which makes it prone to failure with even a small shift in conditions. As such, emerging alternatives are inevitable. Despite having the potential of a revolutionary transformation from traditional education to modern concepts, personalized education faces several challenges. In this article, we discussed these challenges, provided a brief overview of the state-of-the-art research, and proposed some solutions. Table 2 summarizes some of the future research directions.

Acknowledgments

The authors would like to thank the anonymous reviewers and the guest editors for several helpful suggestions.

Authors

Setareh Maghsudi (setareh.maghsudi@uni-tuebingen.de) received her Ph.D. degree (summa cum laude) from the Technical University of Berlin, Berlin, Germany, in 2015, where she is also an adjunct lecturer. She is currently an assistant professor at the University of Tübingen, 72076, Tübingen, Baden-Württemberg, Germany, and a senior project manager at the Fraunhofer Heinrich-Hertz-Institute, 10587, Berlin, Germany. During 2015–2017, she held postdoctoral positions at the University of Manitoba and at Yale University. She has received several competitive fellowships and career awards, including from the German Ministry of Education and Research, the German Research Foundation, and the Japan Society for the Promotion of Science. Her main research interests lie at the intersection of network analysis and optimization, game theory, machine learning, and data science.

Andrew Lan (andrewlan@cs.umass.edu) received his Ph.D. degree in electrical and computer engineering from Rice University. He is currently an assistant professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst, Massachusetts, 01003, USA. He has worked as a postdoctoral research associate at Princeton University. He received a Best Student Paper Award at the 2020 IEEE International Conference on Big Data. He has also organized a series of workshops on machine learning for

Table 2. Some research directions for AI-based personalized education.

Challenge	Description	References
Content production/recommendation	Personalized and profession-oriented production, recommendation, and maintenance of contents	[13], [33]
Evaluation and assessment	Performance comparison in personalized education, testing without information loss, and accreditation	[11], [17]
Lifelong learning	Continuous education and additional qualifications for improvement and advancement in profession	[14]
Incentives	Internal and external motivation for learning, gamification, rewarding, and inducing confidence	[18], [20]
Networking and interaction	Inducing learning networks, forming coalitions for efficient learning, and imitating teacher feedback	[23], [26]
Diversity and fairness	Equal access to a quality online education and avoiding biases in platform development	[27]

education. His research interests lie at the intersection of artificial intelligence, data science, and education.

Jie Xu (jiexu@miami.edu) received his Ph.D. degree in electrical engineering from the University of California, Los Angeles, California, in 2015. He is currently an assistant professor in the Department of Electrical and Computer Engineering at the University of Miami, Miami, Florida, 33146, USA. His main research interests include machine learning and its applications to education and network systems, mobile edge computing, and network security.

Mihaela van der Schaar (mv472@damtp.cam.ac.uk) received her joint bachelor's/master's degree (1996) and her doctoral degree (2001), both in electrical and computer engineering from the Eindhoven University of Technology, The Netherlands. She is the John Humphrey Plummer Professor of Machine Learning, Artificial Intelligence and Medicine at the University of Cambridge, Cambridge, CB2 1TN, U.K. She is also a fellow at the Alan Turing Institute, London, NW1 2DB, U.K., and a Chancellor's Professor at the University of California, Los Angeles, California, 90095, USA. She has received numerous awards, including the Oon Prize on Preventative Medicine from the University of Cambridge (2018), a National Science Foundation CAREER Award (2004), three IBM Faculty Awards, the IBM Exploratory Stream Analytics Innovation Award, the Philips Make a Difference Award, and several Best Paper Awards, including the IEEE Darlington Award. She holds 35 U.S. patents and has contributed to more than 45 international standards. Her research interests include artificial intelligence and operations research for health care and medicine. She is a Fellow of IEEE.

References

[1] Y. Zhong, C. Jiang, W. Xu, and J. Li, "Discourse level factors for sentence deletion in text simplification," in *Proc. AAAI Conf. Artif. Intell.*, Feb. 2020, pp. 9709–9716. doi: 10.1609/aaai.v34i05.6520.

[2] M. Sachan, A. Dubey, E. H. Hovy, T. M. Mitchell, D. Roth, and E. P. Xing, "Discourse ineffectiveness of crowd-sourcing on-demand tutoring from teachers in online learning platforms multimedia: A case study in extracting geometry knowledge from textbooks," *Comput. Linguistics*, vol. 45, no. 4, pp. 627–665, 2020. doi: 10.1162/coli_a_00360.

[3] I. Manickam, A. S. Lan, and R. G. Baraniuk, "Contextual multi-armed bandit algorithms for personalized learning action selection," in *Proc. 2017 IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, pp. 6344–6348. doi: 10.1109/ICASSP.2017.7953377.

[4] C. Magno, "Demonstrating the difference between classical test theory and item response theory using derived test data," *Int. J. Educ. Psychol. Assessment*, vol. 1, pp. 1–11, June 2009.

[5] M. Sharples, "The design of personal mobile technologies for lifelong learning," *Comput. Educ.*, vol. 34, nos. 3–4, pp. 177–193, 2000. doi: 10.1016/S0360-1315(99)00044-5.

[6] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Netw.*, vol. 113, pp. 54–71, May 2019. doi: 10.1016/j.neunet.2019.01.012.

[7] W. Grove and L. Hadsell, *Open Learning Environments, Chapter Incentives and Student Learning*. New York: Springer-Verlag, 2012, pp. 1511–1517.

[8] P. Buckley and E. Doyle, "Gamification and student motivation," *Interactive Learn. Environ.*, vol. 24, no. 6, pp. 1162–1175, 2016. doi: 10.1080/10494820.2014.964263.

[9] P. Vesely, L. Bloom, and J. Sherlock, "Key elements of building online community: Comparing faculty and student perceptions," *MERLOT J. Online Learn. Teach.*, vol. 3, no. 3, pp. 234–246, 2007.

[10] J. Gardner, C. Brooks, and R. Baker, "Evaluating the fairness of predictive student models through slicing analysis," in *Proc. 9th Int. Conf. Learn. Anal. Knowl.*, 2019, pp. 225–234. doi: 10.1145/3303772.3303791.

[11] Z. Wang, A. S. Lan, W. Nie, A. E. Waters, P. J. Grimaldi, and R. G. Baraniuk, "QG-net: A data-driven question generation model for educational content," in *Proc. ACM Conf. Learn. Scale*, June 2018, pp. 1–10. doi: 10.1145/3231644.3231654.

[12] B. Zylich, A. Viola, B. Toggerson, L. Al-Hariri, and A. S. Lan, "Exploring automated question answering methods for teaching assistance," in *Proc. Int. Conf. Artif. Intell. Educ. (AIED)*, July 2020, pp. 610–622. doi: 10.1007/978-3-030-52237-7_49.

[13] A. S. Lan and R. G. Baraniuk, "A contextual bandits framework for personalized learning action selection," in *Proc. Int. Conf. Educ. Data Mining*, June 2016, pp. 424–429.

[14] J. Xu, T. Xing, and M. Van Der Schaar, "Personalized course sequence recommendations," *IEEE Trans. Signal Process.*, vol. 64, no. 20, pp. 5340–5352, 2016. doi: 10.1109/TSP.2016.2595495.

[15] C. Tekin, J. Braun, and M. van der Schaar, "eTutor: Online learning for personalized education," in *Proc. 2015 IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, pp. 5545–5549. doi: 10.1109/ICASSP.2015.7179032.

[16] M. Yudelson, K. Koedinger, and G. Gordon, "Individualized Bayesian knowledge tracing models," in *Proc. Int. Conf. Artif. Intell. Educ.*, July 2013, pp. 171–180. doi: 10.1007/978-3-642-39112-5_18.

[17] A. Ghosh, T. Heffernan, and A. S. Lan, "Context-aware attentive knowledge tracing," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2020, pp. 2330–2339. doi: 10.1145/3394486.3403282.

[18] M. Hartnett, *Motivation in Online Education, chapter The Importance of Motivation in Online Learning*. New York: Springer-Verlag, 2016, pp. 5–32.

[19] S. Maghsudi and M. Davy, "Computational models of human decision-making with application to the internet of everything," *IEEE Wireless Commun.*, early access, Nov. 23, 2020. doi: 10.1109/MWC.001.2000250.

[20] S. D. Levitt, J. A. List, S. Neckermann, and S. Sadoff, "The behavioralist goes to school: Leveraging behavioral economics to improve educational performance," *Amer. Econ. J., Econ, Pol.*, vol. 8, no. 4, pp. 183–219, Nov. 2016. doi: 10.1257/pol.20130358.

[21] E. Deci and R. Ryan, "Motivation, personality, and development within embedded social contexts: An overview of self-determination theory," in *Oxford Handbook Human Motivation*, R. M. Ryan, Ed. London, U.K.: Oxford Univ. Press, 2012, pp. 85–107.

[22] A. Bandura, *Self-Efficacy: The Exercise of Control*. San Francisco: Freeman, 1997.

[23] Y. Xiao, F. Dörfler, and M. Van Der Schaar, "Incentive design in peer review: Rating and repeated endogenous matching," *IEEE Trans. Netw. Sci. Eng.*, vol. 6, no. 4, pp. 898–908, 2018. doi: 10.1109/TNSE.2018.2877578.

[24] S. Maghsudi and M. van der Schaar, "Distributed task management in cyber-physical systems: How to cooperate under uncertainty?" *IEEE Trans. Cogn. Commun. Netw.*, vol. 5, no. 1, pp. 165–180, 2019. doi: 10.1109/TCCN.2018.2888970.

[25] Y. Zhang and M. van der Schaar, "Strategic networks: Information dissemination and link formation among self-interested agents," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 6, pp. 1115–1123, 2013. doi: 10.1109/JSAC.2013.130613.

[26] A. S. Lan, J. Spencer, Z. Chen, C. Brinton, and M. Chiang, "Personalized thread recommendation for MOOC discussion forums," in *Proc. Eur. Conf. Mach. Learn. Principle Knowl. Discov. Databases*, Sept. 2018, pp. 725–740. doi: 10.1007/978-3-030-10928-8_43.

[27] J. Reich and M. Ito, *From Good Intentions to Real Outcomes: Equity by Design in Learning Technologies*. Irvine, CA: Digital Media and Learning Research Hub, 2017.

[28] S. Doroudi and E. Brunskill, "Fairer but not fair enough on the equitability of knowledge tracing," in *Proc. 9th Int. Conf. Learn. Anal. Knowl.*, 2019, pp. 335–339. doi: 10.1145/3303772.3303838.

[29] E. Garcia and E. Weiss, "COVID-19 and student performance, equity, and U.S. education policy: Lessons from pre-pandemic research to inform relief, recovery, and rebuilding," Economic Policy Institute, Washington, D.C., 2020. [Online]. Available: https://files.epi.org/pdf/205622.pdf

[30] G. Di Pietro, F. Biagi, P. D. Mota Da Costa, Z. Karpinski, and J. Mazza, "The likely impact of COVID-19 on education: Reflections based on the existing literature and recent international datasets," Publications Office of the European Union, Luxembourg, Tech. Rep. EUR 30275 EN, 2020. doi:10.2760/126686, JRC121071.

[31] A. Dixit, U. S. Shanthamallu, A. Spanias, V. Berisha, and M. Banavar, "Online machine learning experiments in HTML5," in *Proc. IEEE Frontiers Educ. Conf.*, 2018, pp. 1–5.

[32] U. S. Shanthamallu, S. Rao, A. Dixit, V. S. Narayanaswamy, J. Fan, and A. Spanias, "Introducing machine learning in undergraduate DSP classes," in *Proc. Int. Conf. Acoustics, Speech Signal Process.*, 2019, pp. 7655–7659.

[33] N. J. Cepeda, N. Coburn, D. Rohrer, J. T. Wixted, M. C. Mozer, and H. Pashler, "Optimizing distributed practice: Theoretical analysis and practical implications," *Exp. Psychol.*, vol. 56, no. 4, pp. 236–246, 2009. doi: 10.1027/1618-3169.56.4.236.

SP

Clive Cheong Took, Stephen R. Alty,
Anush Yardim, and David M. Howard

Creativity First, Science Follows

Lessons in digital signal processing education

©SHUTTERSTOCK.COM/LUCKYSTEP

Digital signal processing (DSP) education has traditionally employed more demanding mathematics than most topics found among courses in electrical/electronic/computer engineering. In some cases, the technical challenges posed by some courses have made it difficult for students to complete those courses successfully. Here, we advocate for creativity to be nurtured in the first place, after which the science will flow naturally. To foster creativity, our pedagogical approach includes a variety of solutions incorporating exploratory exercises, open-ended multidisciplinary coursework, blended lecture–laboratory sessions, and a colorful working environment. We firmly believe that creativity is the way forward. Student feedback supports our approach.

Introduction

Royal Holloway, University of London (RHUL) took the decision to open a new Department of Electronic Engineering, with its first cohort of students enrolling in 2017 [1]. This gave us a welcome opportunity to develop our electronic engineering course, and subsequently computer systems engineering, around a strong strand of DSP, adopting an experimental perspective free from legacy commitments that will foster the creativity so readily achievable in modern-day DSP practice.

The common trend indicates that research and academia tend to focus on the theoretical solution (by proposing new mathematical models and algorithms), whereas industry spends more time on solving the problem (by understanding the data) [2]. We are of the opinion that these two elements need to be balanced out. To strike such a balance, we aim to assess student understanding and performance proficiency by a mix of approaches that incorporate computer-supported simulations, data exploration, and traditional hand calculations.

Related works

Creativity stems from divergent thinking (i.e., the generation of ideas) rather than convergent thinking (i.e., analysis and evaluation) [3]. Analysis and mathematical rigor are part and parcel of signal processing, such as bounded-input, bounded-output (BIBO) stability for filters, convergence, and steady-state analyses for

Digital Object Identifier 10.1109/MSP.2021.3058459
Date of current version: 28 April 2021

adaptive filters. As such, it is typically difficult to inculcate creativity in DSP-oriented modules. However, there are many research efforts in other engineering disciplines that help creative thinking/learning. As there is not one kind of education that fits all [4], a set of good practices and their corresponding learning outcomes for creativity are as follows:

- exploratory exercises allowing students to investigate/explore new ideas or concepts or models on their own [5]
- open-ended problems that can be solved in a multitude of ways for students to think independently [6]
- learning opportunities for students to take independent responsibilities or initiatives [6]
- collaborative work for students to brainstorm and generate ideas [7]
- multidisciplinary approaches whereby students borrow principles from other engineering domains to solve problems [8].

Most importantly, the value of instructors believing in creativity was particularly highlighted in [9]. While believing in creativity is not a method for teaching creativity, it is a decisive factor in teaching it [7], [8]. However, the research did not consider the particularities of DSP education. To this end, our course development focuses on practical skill sets and promotes creativity by considering all of these pedagogical sets of good practices. We have designed three final-year DSP modules:

- Digital Signal Processing Design (EE3010), which provides a grounding in DSP practicalities
- Fundamentals of Biomedical Engineering (EE3060), which gives the opportunity for students to explore biomedical signals and systems
- Voice Technologies (EE3050), which offers diverse applications of speech signal processing ranging from voice cloning to voice forensics.

The rationale for these DSP modules stems from the demands from industry for qualified graduates in such enabling technology for communications, sensors and instrumentation, medical applications, very large-scale integration, avionics, audio industries, radar, and many other key sectors. To put into context how these three DSP modules fit into the undergraduate study, Table 1 shows our overall program; the pathway highlighted in yellow illustrate the DSP theme. For example, the prerequisite for EE3010 is the course Signals, Systems, and Communications (year 2), and that of EE3050 is EE3010. However, there is also some degree of interdependence between the different pathways. For instance, the Fourier series taught in Mathematics for Engineers 2 (the general engineering theme) could be used to model the periodicity of electrocardiograms (ECGs) in Fundamentals of Biomedical Engineering (the DSP theme). Students exploit their background of electronic circuit designs from the blue pathway and embedded systems from the purple pathway and their DSP knowledge from the yellow pathway in EE3060.

One of the innovations of the Digital Signal Processing Design (EE3010) module is to adopt a blended approach by delivering the lecture session in a laboratory environment. Thus, the students move between theoretical concepts and immediate practical illustrations—a teaching strategy to make the mathematical content more engaging. On the other hand, the other courses, EE3050 and EE3060, do not take this blended approach as they are inherently application oriented (where students can more readily contextualize their usefulness). At the end of course EE3010, the students are expected to be able to

- examine the scientific principles underpinning practical signal processing and apply the knowledge gained from the major aspects of DSP to solve problems efficiently
- apply a modeling approach to engineering problems to appreciate the application of relevant technologies in signal processing
- design systems using effective software instrumentation tools that facilitate rapid proof of concept.

Our DSP analysis emphasis is on the verification of DSP time- and transform-domain relationships closely supported by illustrative simulation experiments relying on MATLAB, Simulink, and DSP_Speedster [10], the latter allowing for a virtual instrumentation environment that piggybacks on the previous two. We are especially keen on Simulink models because they convey dynamic DSP scenarios, which are nearly as effective as benchtop instrumentation exercises while also furnishing flexibility through easy program/parameter changes. This is aided by simultaneously appealing to students'

Table 1. Creativity is inherently embedded into our undergraduate program, e.g., creative team project 1. The color-coded legends show the themes of courses, i.e., yellow (DSP), blue (electronic circuit design), orange (computer science), gray (general engineering), pink (project), and green (sustainable and power engineering). All courses are mandatory except for those labeled with (O).

Year 1		Year 2		Year 3	
Term 1	Term 2	Term 1	Term 2	Term 1	Term 2
Embedded systems creative team project 1		Embedded systems creative team project 2		Final year project	
Electronic Circuits and Components	Communications Engineering	Signals, Systems, and Communications	Electronic Materials and Devices	Introduction to Project Management	Digital Systems Design
Programming in C++	Internet Services	Analog Electronic Systems	Professional and Sustainable Engineering	Digital Signal Processing Design	Power Systems (O)
Mathematics for Engineers 1	Mathematics for Engineers 2	Software Engineering	Control Engineering	Fundamentals of Biomedical Engineering (O)	Advanced Communication Systems (O)
(O) = optional				Information Security (O)	Voice Technologies (O)

visual and auditory senses through observing soft-instrumentation scopes and listening to outputs. Such time-varying/adaptive model usage far surpasses the motivational impacts provided by static analysis through pencil and paper exercises or even MATLAB-powered numerical tabulated solutions.

The DSP_Speedster lab kit environment is especially helpful when dynamic evolutionary situations (such as time-varying filtering or adaptive processing) require monitoring and control. “Snapshotting” through static MATLAB computations often falls short of providing the required operations insight, where benchtop instrumentation could excel. Yet this is where Simulink modeling delivers a computationally attractive virtual instrumentation alternative. Simulink/DSP_Speedster modeling is quick and natural, establishing a bridge to later, more extensive design and prototype refinement activities.

We fully appreciate and acknowledge the necessity for tight coupling between theory and practical realizations. All experimental work is motivated and focused as a follow up to the preliminary theoretical background, which is always pivotal in a field like DSP, having such exceptional and intrinsic alignment with mathematical underpinnings. Our approach is to blend analysis and preparatory hand calculations with MATLAB-supported analysis and plots of expected performance. This is what we feel is the static phase of the student’s journey toward creative design. The dynamic phase comprises the instrumentation and exercising of the operational aspects of solution implementations. This stage typically involves benchtop activity with extensive laboratory equipment alongside Simulink modeling and performance assessments. Finally, the assimilation of findings and reflection on results informs a fresh wave of experimentation and refinement.

The two examples that follow are typical scenarios that our students experience in the third year of EE3010 in the DSP design module. There are similar experiments that second-year students undertake as well to bolster their understanding of modulation and modern communication trends, such as software-defined radio. The learning outcomes for creativity of these two examples encourage students to explore new models on their own through the development and evaluation of those models and take independent responsibility and initiative.

Example 1: BIBO stability experiment

Students are expected to build the model given in Figure 1 and investigate the behavior of this system (running at 20,000 samples/s) under dynamic conditions, for example, by first varying the feedback parameter in the slider gain block. The input is a periodic impulse train with a period of 300 samples. Here, “del” is a differencer, and “TPMA” is a two-point moving averager with impulse responses $\{1, -1\}$ and $\{0.5, 0.5\}$, respectively.

Each student builds the model depicted in Figure 1 and commences experimentation. The questions they need to answer are whether the system is BIBO-stable [11, p. 24] as the model is initially specified and which range of the feedback parameter values will make the system become stable/unstable. Consequently, students are invited to explore and be curious about this system at first. So far, no equations are needed. They realize early on how a system that they built can become unstable very quickly under certain conditions. Normally, this is not experienced by students in most courses; they are given either a stable or an unstable system, and once they find out that the poles are outside of the unit circle, they declare instability—with little appreciation of the journey that led to it.

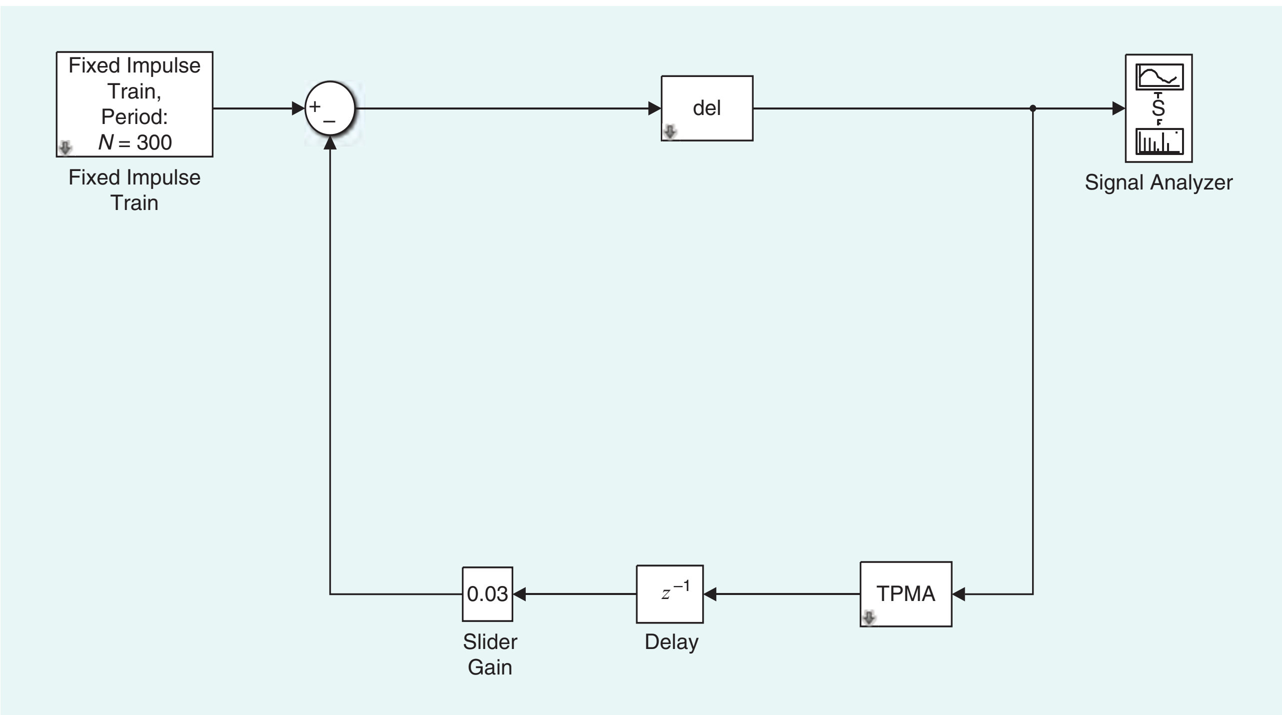

FIGURE 1. A stability investigation of a feedback system.

In this scenario, students have to explain why the system behaves this way by obtaining the z-domain transfer function and analyzing its pole–zero pattern (PZP) after they encounter the unbounded nature of the output. This gives the motivation for analytical exploration; students see the "what" first, and now they have to answer the "why." Extensions to the investigation involve swapping the "del" and "TPMA" blocks and other configurations in the feedforward and feedback paths to investigate the model behavior. Within DSP_Speedster, they have access to many other blocks, such as more exotic filters, that can be introduced. Immediately, the investigation becomes not only individualized (thus avoiding collusion) but also enables students to think creatively to design an overall BIBO-stable system. A further extension is to design a compensator to guarantee an overall set of specifications by exploring cascade and parallel compensator configurations.

All results, models, and plots, including impulse response behavior, spectral gain, and PZP, are then submitted online together with a brief report. A typical plot of the unstable response is given in Figure 2.

This style of experimental activity has provided an excellent framework for practical MATLAB/Simulink-based tests and examinations during the academic term. Feedback from past students indicated a strong preference for these rather than the classical pen and paper examination. After all, this is very much how we engineers operate in the real world—by being curious and creative with the support of well-crafted tools to solve open-ended problems.

Example 2: Adaptive notch filtering

This experiment primarily aims at eliminating additive tonal interference from a background random white process using an adaptive notch filter. Meanwhile, an alternative learning viewpoint is that (as well as achieving tone removal) this system moreover furnishes a useful frequency estimation capability for noisy tone-hopping situations [12], [13]. Students design and operate the system depicted in Figure 3, where they manipulate the sinusoid frequency, the convergence factor (mu), and the signal-to-interference ratio.

The students adjust the parameters in the signaling environment such that the system is able to satisfactorily home onto and suppress the contaminating tone. The "Seepig Filter Analyzer," which is not available in mainstream Simulink [10] but is part of DSP_Speedster, enables the instantaneous tracking and visualization of any changes in the PZP, impulse response, and gain of the filter. While it is not highlighted here, other filter characteristics, such as group delay, phase, phase delay, zero-phase gain, total impulse response energy, average delay, and impulse response center of gravity, are readily available for dynamic measurement and display. No gradient-search adaptive notch filter like the DSP_Speedster block seems to be furnished in standard Simulink or in the DSP toolbox. The entire search algorithm is realized in elemental Simulink blocks, and its detailed action can be viewed by students.

The green plots in Figure 4 are time and spectrum plots before the experiment starts, and the red plots indicate when the tuning error has converged to zero.

All of the exercises are totally paperless; students receive their instructions by opening the blue "info" box in the top right corner of the model (Figure 3). Of course, it must be noted that students also experience the sound of the tone before it is extinguished as well as the narrowband sweeping chirp as the notch filter steers toward the tone's spectral location. So, they see, hear, and absorb—they are being creative in their exploration.

For conciseness, only isolated specimen experiments are presented. Students have undertaken MATLAB-based (which includes DSP_Speedster) practical exams, and feedback is strongly in favor of the approach that we have reported here.

In the Fundamentals of Biomedical Engineering (EE3060) course, the students learn about biosignal processing techniques (e.g., time–frequency analysis) as well as the particularities of the biosignals. However, biosignal processing goes beyond these two elementary know-hows. To have a more holistic view on biosignal processing, students are expected to exploit their practical knowledge gained from nonsignal processing courses (e.g., analog electronics and embedded systems) to solve signal processing

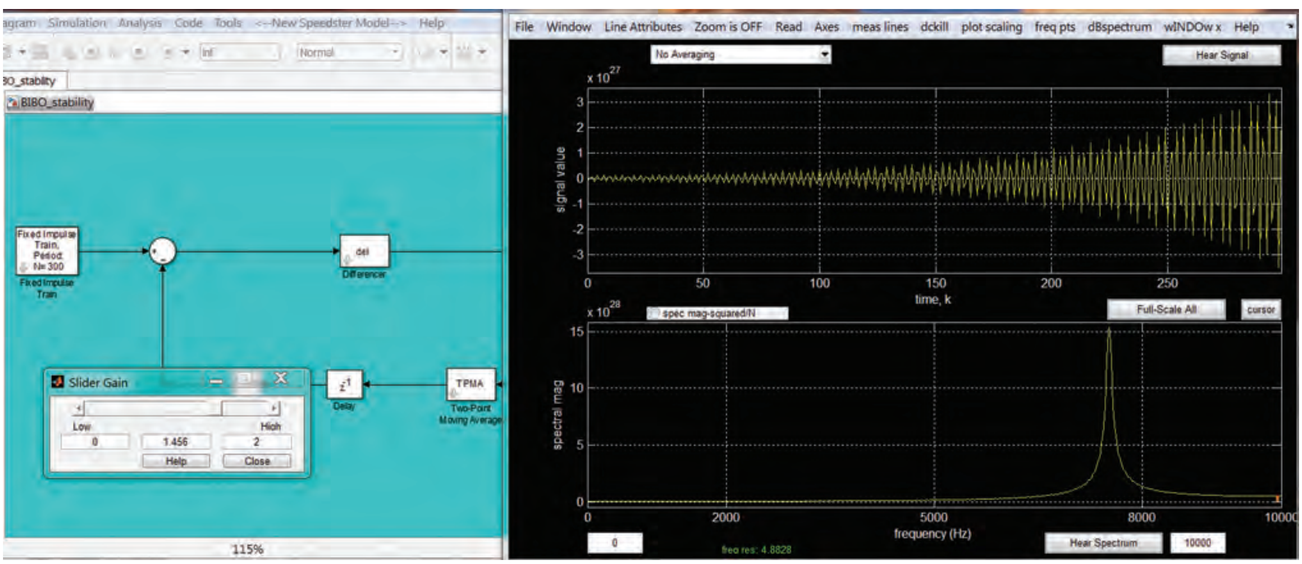

FIGURE 2. A snapshot showing BIBO instability.

problems. As such, by the end of this course, the students should be able to work in teams to do the following:

- Address and analyze problem-driven (instead of theory-driven) DSP tasks with no unique solution and with no constraints except financial constraint. This type of open-ended assignment gives the opportunity for students to be creative. An example of such a problem-based assignment is to develop a smart system to detect drowsiness and alert the individual, as illustrated in Figure 5.
- Build a complete DSP system from start to finish. This involves designing the data acquisition system (i.e., circuit analysis and implementation), investigating the appropriate biosignal processing algorithm to undertake the real-time analysis (e.g., frequency analysis), developing a smart system (e.g., coding an embedded system) to actuate on the results of the DSP data analysis, and manufacturing 3D objects to improve the esthetics of their DSP product.

To achieve these outcomes, students learn about:

- *Biodata exploration*: To be able to exploit the properties of biomedical signals, the students learn about the particularities of electroencephalogram (EEG) for the brain, the ECG and the photoplethysmogram for the heart, and the electromyogram

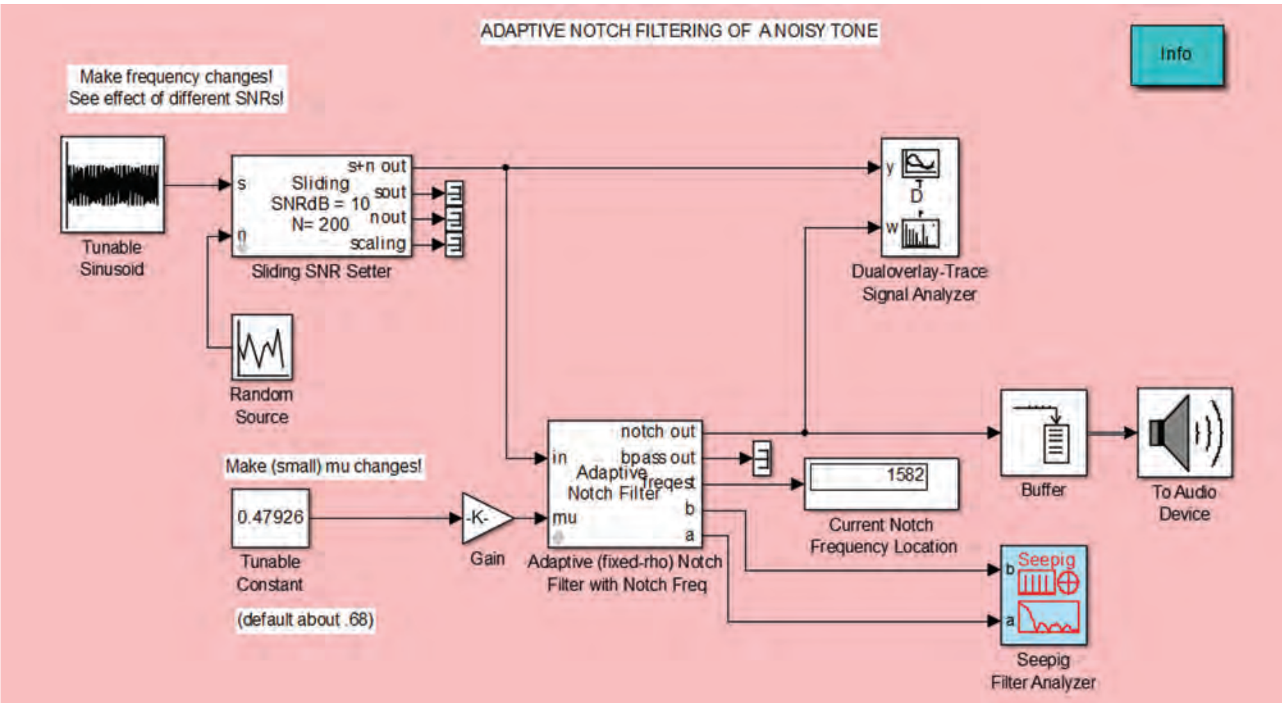

FIGURE 3. An adaptive notch filtering model. SNR: signal-to-noise ratio.

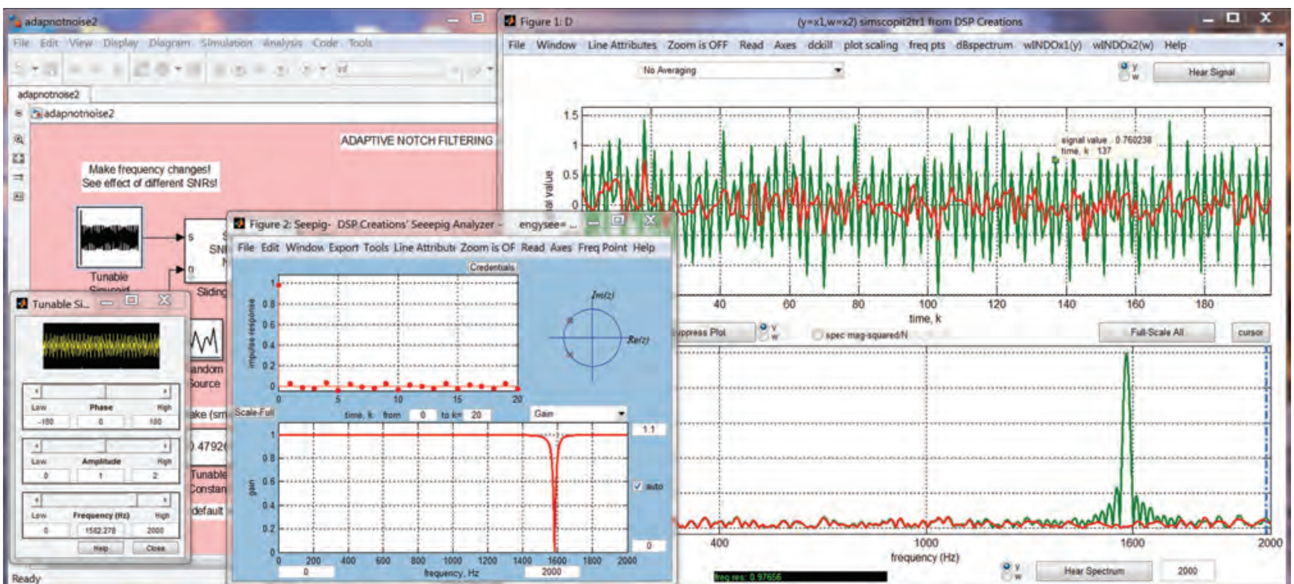

FIGURE 4. A notch filter adaptation for a noisy sinusoid.

(EMG) for muscles. Moreover, the students also learn about medical imaging and appreciate the difference between anatomical and functional imaging.

- *Biosignal acquisition and instrumentation*: The students design circuits for two main purposes—filtering and amplification. Examples of filtering applications include notch filtering to remove a 50-Hz power line, high-pass filtering to remove motion artifacts, and low-pass filtering to alleviate unwanted noise outside the frequency range of the biomedical signal. Amplification of those biosignals is achieved via both differential amplifiers and instrumentation amplifiers. Differential amplification is particularly useful for signals such as ECG, which is measured using the potential difference between the left and right side (polarity) of the body. This difference also helps in terms of denoising as the instrumentation amplifier inherently cancels the common-mode noise.
- *Biosignal processing and learning methods:* Students learn about windowing and its effect on the spectral properties of the windowed data, time–frequency analysis, and signal-dependent methods, such as for the detection of the QRS complex in ECG and the segmentation of EMG in terms of muscle contraction and relaxation. Students also gain fundamental knowledge in terms of real-time learning algorithms (such as the perceptron). Thereafter, the students can take advantage of these learning methods to make automated decisions in their bio-DSP system.

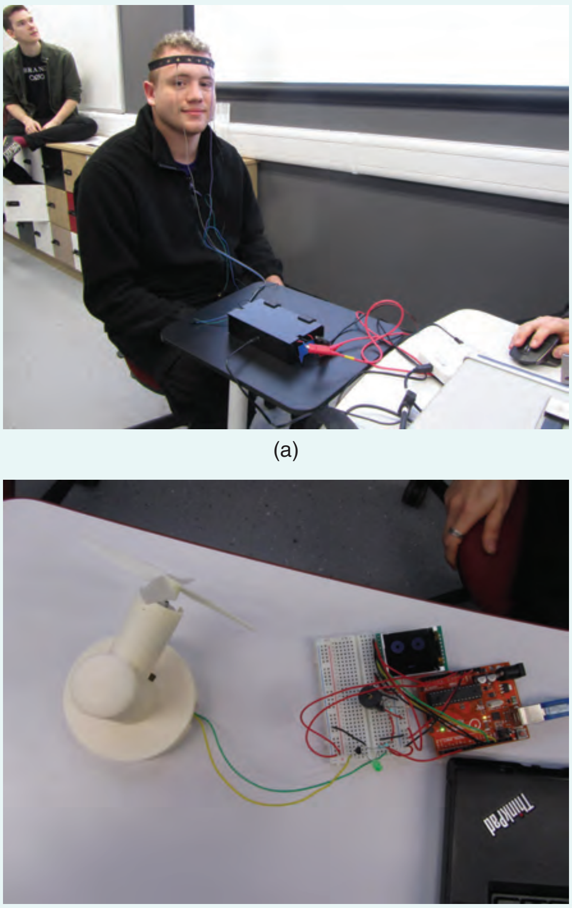

(a)

(b)

FIGURE 5. (a) An EEG sensing system in a 3D printed enclosure to detect drowsiness. (b) An actuating system to provide response to a drowsiness state via a 3D printed fan, LCD display, and melody player (buzzer).

For the group coursework on building a complete bio-DSP system, the learning outcomes for creative learning are to enable students to 1) brainstorm and generate ideas, 2) borrow principles from other engineering modules to solve problems, and 3) take independent initiative.

Example 3 gives us a flavor of the type of experiment carried out in Fundamentals of Biomedical Engineering. In this exercise, the students are not only exposed to hardware experimentation to capture the biosignal but also take a "white-box" approach to generate such biosignals synthetically. This approach reflects one of the strengths of DSP, giving us control over the design parameters. And its learning outcomes for creative learning are to empower students to think independently and take independent responsibility.

Example 3: QRS analysis and synthesis for ECG

In this experiment, the students "learn by doing" ECG analysis. The common approach is to acquire ECG signals and then undertake peak-to-peak analysis (such as between two R peaks). However, the students are then asked additionally to generate their own ECG signals synthetically. The value of such an experiment is that 1) it gives the opportunity to the student to be creative (e.g., the periodicity of ECG cycle can be achieved in various ways, such as Fourier series of an ECG cycle or looping an ECG cycle); 2) students appreciate the variability of biomedical signals; 3) it reinforces the learning of the student "by doing" rather than by memorization; and 4) the synthetic data can be used as "controlled data" as too often students overlook the importance of controlled experiments in biomedical analysis. An example of a real-world ECG signal and its corresponding synthetic version are illustrated in Figure 6.

The aim of the Voice Technologies (EE3050) course is to provide students with advanced knowledge of voice production, synthesis, recognition, and processing in the context of present-day and future engineering systems that make use of a voice input or output. At the end of this course, the student should be able to

- examine the engineering principles and techniques necessary to understand and analyze how voice can be created or recorded, processed, stored, and delivered to the user
- apply a holistic approach to voice synthesis, recognition, and processing through the application of the relevant technologies
- show the context in which engineering knowledge can be applied to voice synthesis, recognition, and processing
- extract and evaluate pertinent data and apply engineering analysis techniques in the solution of unfamiliar problems.

As such, the students have the opportunity to learn about 1) human voice production for speech and singing; 2) electronic synthesis of human speech and singing in terms of the sound source and sound modifiers to create synthetic voice signals; 3) signal processing techniques used, for example, to track vocal pathologies, monitor the changes in vocal skills during voice trauma recovery speech therapy or the development in vocal skills during acting or singing voice training, enhance voice quality, remove background noise, and improve perceived voice quality; 4) the design of hearing aids including cochlear implants; and 5) techniques used for automatic speech recognition, such as Apple's Siri system.

The innovation in this course is to focus on the technologies rather than the mathematical models in speech processing. Students tend to learn and understand more about a subject if they can appreciate its application and therefore find it useful. Thus, the applicability of these speech technologies is illustrated in synthetic voice generation, hearing aid design, clinical and research voice monitoring systems, the impact on perceived voice quality, and an overall understanding of the spoken message. On the other hand, a traditional speech processing course would typically focus on mathematical models, such as the Levin–Durbin recursion for linear speech prediction or the derivation of the transfer function of the vocal tract with poles and zeros. We do not take this traditional approach. As such, our laboratory sessions are based on voice technologies widely used in the speech community, such as Audacity [14] and Praat [16].

Audacity [14] is an open source multitrack editor and recorder for audio recordings. Students make a recording of their own voice in the first week, which includes isolated words, counting forward and backward from 0 to 20, and reading a section of the phonetic read passage "Arthur the rat" [15]. They then prepare their signals for analysis using Audacity, thereby having a hands-on learning experience directed toward the transmitted signal itself rather than any underlying mathematical models.

Praat [16] is a free tool for phonetics research that enables the students to do speech capture, manipulation, waveform, and spectral analysis as well as formant and articulatory synthesis. In our laboratory sessions, Praat enables students to study multiple items.

- *Time-domain analysis*: The students isolate individual spoken sounds and measure their durations where appropriate while gaining an understanding of the dynamic nature of running speech and transitions between phonemes.
- *Frequency-domain analysis*: The students explore the formant structures of different vowels, with special exercises relating to the effect on the output of varying the analysis filter bandwidth in the context of wideband and narrow-band spectrograms, particularly in the context of the dynamic nature of formant transitions in diphthongs and the spectral nature of consonants during running speech.
- *Time- and/or frequency-domain analysis*: Fundamental frequency estimation is explored in the time and/or frequency domain in the context of a hands-on experience of the advantages and disadvantages of each approach in the context of human speech as well as the acoustic analysis of "connected" speech, such as the acoustic analysis of syllables and the analysis of a word in different contexts.
- *Linear predictive coding (LPC)*: The students investigate the frequency response of the vocal tract and that of the sound source through LPC and its application in telephony. Having used LPC to code and decode a speech signal, they attempt to resynthesize speech, having replaced the larynx input with nonspeech sources, such as music for fun, along the lines of Sparky's Magic Piano [17].
- *Voice cloning*: Students are able to explore time- and frequency-domain differences in the speech of different speakers in the context of why they sound different and yet the spoken

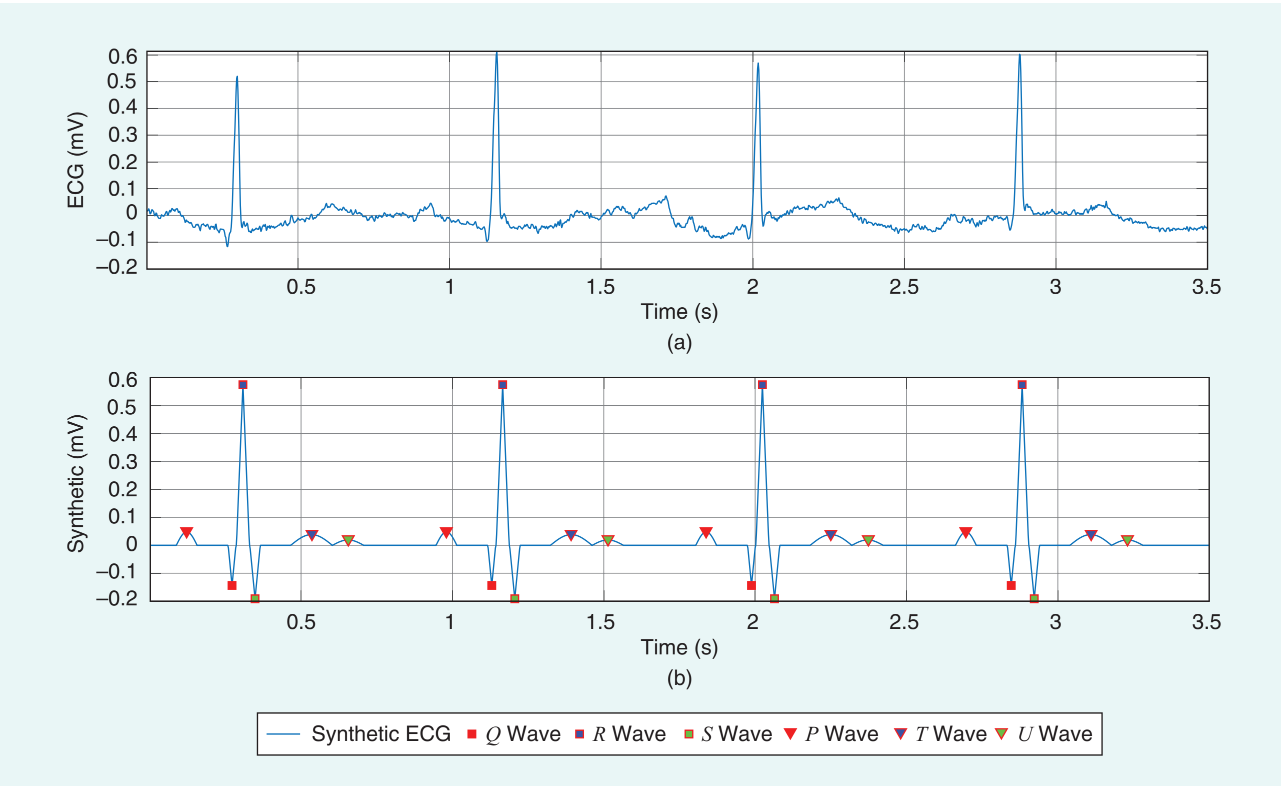

FIGURE 6. An ECG signal and its corresponding synthetic counterpart. This exercise encourages the students to appreciate heart variability and reflect more carefully on the different ECG wave components, especially those that are not so visible in the ECG signal.

message can still be understood, firstly through the synthesis and analysis of different vowels, and then through running speech generation using the CereProc [18] online speech synthesis system.
- *Hearing loss*: Having learned the principles of human hearing, students explore the frequency-domain nature of their own hearing (via headphones and being aware of the local acoustic noise) using a simple audiometer implemented in Pure data or Pd—an open source graphical programming audio creation and manipulation system [19]. In addition, having explored noise-induced hearing loss (NIHL) introduced in a lecture, students perceptually investigate which speech sounds should be affected adversely and then test their hypotheses by exploiting notch filters in their laboratory session to mimic the spectral (not the signal level as this would pose a direct health and safety threat) effects of NIHL and confirm (or otherwise) their predictions.

Example 4: Forensic analysis of curious sounds

This experiment allows students to explore "curious" voice sounds set in a context of forensic audio comparisons that have been discussed in the associated lecture, including an original voice of a person and other related voices, such as
- his voice after inhaling helium saying the same words
- his voice mimicked by a professional
- formants as sine waves [20]
- a voice from a talking elephant and a seal
- the Laurel and Yanny illusion voice [21].

Spectrographic analysis is linked to the basis of how the hearing system works, and some creative lateral thinking is encouraged through the consideration of nonspeech sounds. Students are asked to analyze sounds such as those depicted in Figure 7 and to think about how they have been created. As such, the learning outcomes for creative learning were to enable students to think independently and take independent responsibility or initiative.

To focus on the practical skills acquired by the students, the main assessment for voice technologies was a 2-h practical examination that ended up being taken remotely by students who had left campus due to the COVID-19 pandemic. Each student was given a different (to avoid any direct collusion of measured values) spoken version (16-bit, 44,100-Hz sampled mono .wav file) of the sentence "She said to her friend, can I go out tonight to see the opera with you?" to be phonetically transcribed and analyzed in terms of fundamental frequency statistics, formant frequencies of a selection of vowels to enable a link to be made with their tongue position within the vocal tract, the nature of frication energy for a few fricatives, the acoustic similarities and differences in the three "n" sounds in the sentence, and the acoustic nature of sentence stresses.

The assessment was designed as a creative exercise, where students had the freedom to choose any DSP analyses. In their solution, the students carried out those analyses, from which they were

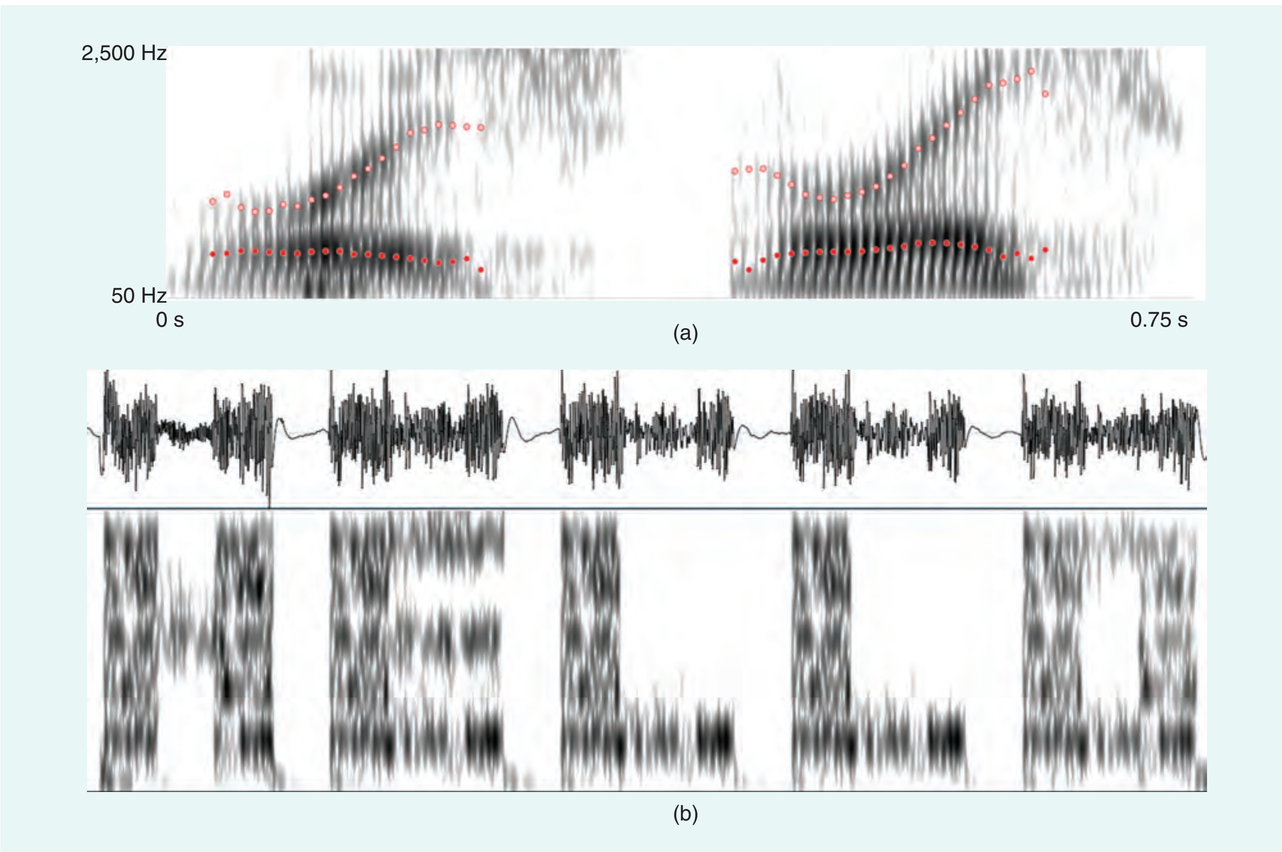

FIGURE 7. (a) A wideband spectrogram analysis of the word "voice" spoken in air (left) and in helium (right) with the first two formants indicated by the red dots, showing the upward shift in formant frequencies due to helium. (b) A time waveform (upper) and wide-band spectrogram (lower) of a "curious" sound that is left as a creative thinking exercise for the student to explain how it was created.

expected to appreciate the acoustic variation in speech output from an individual speaker—something that would be highly relevant for speech recognition, transmission, synthesis, and storage.

Unlike traditional DSP courses that focus on communication-based problems, these three courses facilitate student learning in relation to their everyday life activities and experiences, e.g., their usefulness in wearables, healthcare tracking, and forensic technologies. However, these DSP courses alone are not adequate to foster creativity and the learning process, which leads us to the next factor in our endeavor to enhance learning—the environment.

The environment

To maximize student engagement, a lecture theater [Figure 8(a)] was jointly designed by students and academics to

- allow students to be seated in teams so that peer learning can be facilitated
- enable the lecturer to roam around the class (including between each row of seats) so that students are not in their comfort zone, and the lecturer is not confined to the front of the class
- use different color lighting to gauge the mood of the students. It has been reported in chromotherapy [22] that red light can stimulate the body and mind and increase circulation (e.g., during important parts of the lecture), whereas blue light is believed to soothe illnesses and treat pain (e.g., during breaks within a lecture).

Our lecture theaters are also equipped with the Panopto video platform [23], which captures our lecture sessions. This allows our students to catch up with missed lectures or even revisit the lecture when things start "clicking."

For brainstorming sessions where creativity is key, we have adopted the Google approach: we have our creative thinking room [Figure 8(b)] that provides ample and colorful space for students to have lightbulb moments with adaptable furniture and screen displays for discussions.

In the lecture theaters, everyone is the same. All academics are seated in an open-plan office. This inclusivity does not stop with academics. All students of electronic engineering have access to the open-plan office. Our open-door (office) policy encourages our students to engage with academics with impromptu discussions—when creativity comes to light and science follows.

Feedback on the courses

To evaluate the impact of our innovative approach in teaching, Tables 2 and 3 summarize two surveys from our students. The first survey in Table 2 was undertaken externally by an agency, Ipsos MORI; this survey is known as the National Student Survey and therefore provides us a benchmark against other universities in the United Kingdom [24]. However, the first survey does not focus specifically on creativity in teaching. To this end, a second survey on creativity (Table 3) was carried out to investigate the impact of our teaching on creativity. The first survey interviewed our first cohort of graduates (13 students), whereas the second survey was based on our current cohort (30 students).

Discussion

Although not all questions in Table 2 are directly relevant to our approach on creativity, the survey does offer a useful benchmark at the national level. Relevant questions are highlighted in bold in Table 2, whereas all of the questions in Table 3 focus on creative and practical learning.

Student engagement initiatives

More than eight out of 10 students found that the staff made the subject interesting. Perhaps this is due to our ongoing effort to contextualize the theories with applications, e.g., notch adaptive filtering in audio applications or ECGs in wearables [25]. We have always endeavored to make our course intellectually stimulating by exposing students to open-ended problems or exploratory exercises. In fact, nine out of 10 students agreed. Likewise, the statistics for the question on learning opportunities to explore ideas in depth in Table 2 corroborate with those of question 5 in Table 3. Another contributing factor (not discussed in this article) that engages students while promoting their independent and creative thinking is research projects, as found in [26].

Practical-oriented teaching

Our project-led coursework encouraged our students to bring information together from different topics. For instance, they had to

(a)

(b)

FIGURE 8. (a) The lecture theater designed jointly by students and academics. (b) The creative thinking room with funky furniture whose colors were inspired to reflect the resistor color codes.

apply concepts from circuits and embedded systems to solve a biomedical signal processing problem in EE3060. Our practical approach to teaching has been successful, with 92% of our students acknowledging that they applied what they have learned. It is not a surprise, therefore, that most students prefer practical-based examinations (see question 1 in Table 3). The importance of the practical element in DSP education has already been highlighted [27].

Table 2. The feedback on our teaching from our first graduate cohort in 2020. More details available from the National Student Survey [24]. Questions in bold font are relevant to our innovation in education, i.e., creativity and openness, including easy accessibility to staff.

Questionnaire	Actual Value	Sector Average in the United Kingdom
Teaching on my course		
Staff are good at explaining things.	83%	84%
Staff have made the subject interesting.	**83%**	**75%**
The course is intellectually stimulating.	**92%**	**86%**
Learning opportunities		
My course has provided me with opportunities to explore ideas or concepts in depth.	**83%**	**78%**
My course has provided me with opportunities to bring information and ideas together from different topics.	**83%**	**82%**
My course has provided me with opportunities to apply what I have learned.	**92%**	**78%**
Assessment and feedback		
The criteria used in marking have been made clear in advance.	83%	67%
Marking and assessment has been fair.	83%	74%
Feedback on my work has been timely.	92%	61%
I have received helpful comments on my work.	75%	64%
Academic support		
I have been able to contact staff when I needed to.	**92%**	**86%**
I have received sufficient advice and guidance in relation to my course.	92%	76%
Good advice was available when I needed to make study choices on my course.	83%	71%
Organization and management		
The course is well organized and is running smoothly.	92%	65%
The timetable works efficiently for me.	92%	78%
Any changes in the course or teaching have been communicated effectively.	83%	75%
Learning resources		
The IT resources and facilities provided have supported my learning well.	83%	84%
The library resources (e.g., books, online services, and learning spaces) have supported my learning well.	75%	84%
I have been able to access course-specific resources (e.g., equipment, facilities, software, and collections) when I needed to.	83%	87%
Learning community		
I feel part of a community of staff and students.	83%	67%
I have had the right opportunities to work with other students as part of my course.	**100%**	**89%**

Environment

Our open-door policy also facilitated students getting prompt feedback on their work as well as academic support in general. Table 2 confirms that this is the case. Students believe creativity is crucial in engineering, and they have been encouraged to be creative in their work (e.g., questions 3 and 8 in Table 3). On the other hand, the students did not value the working environment as much as other factors. Only six out of 10 students believed that the creative thinking room was helpful in their study. Although the impact of this factor is not as apparent as the others, it is in the creative thinking room where the students would typically brainstorm. We tend to value the product design rather than the product process, which might explain the lower statistics for the environment [9]. In fact, it was found by several researchers that color and furniture play an important role in creativity [28], [29]. Our creative thinking room [illustrated in Figure 8(b)] provides our students such an environment.

Conclusions

Divergent thinking leads to creativity, yet we are trained to focus on convergent thinking when we emphasize evaluation and analysis [3]. There is not just a single kind of education that can teach creativity. As such, we have adopted a variety of good practices to encourage our students to be creative. These include open-ended problems, exploratory laboratory exercises, project-based coursework that requires multidisciplinary and teamwork skills, and a creative working environment. Student feedback confirms that creativity is an important aspect of engineering. We hope that this article encourages educators to take more risks and embed creativity in

Table 3. The student survey specifically addressing creativity and practical examinations.

Questionnaire	Yes	No	Neither
Practical-based examinations			
Q1. Do you feel that practical-based examinations are more appropriate than traditional paper-based examinations to assess your technical knowledge?	83%	7%	10%
Q2. Do you feel that practical-based examinations are more appropriate for students with disabilities than traditional paper-based examinations?	57%	17%	26%
Creativity			
Q3. Do you think creativity is an important aspect of engineering?	100%	0%	0%
Q4. Do you feel your creativity is stretched more by open-ended coursework than coursework with unique solutions?	83%	7%	10%
Q5. Has the open-ended coursework motivated you to research materials beyond the materials available for the module?	77%	17%	6%
Q6. Was the creative thinking room helpful in your studies?	60%	13%	27%
Q7. Have learning opportunities that fostered your creativity at RHUL consolidated your independent thinking?	77%	10%	13%
Q8. Do you feel that you have been encouraged to be more creative by your study at RHUL?	70%	10%	20%

their DSP teaching. We find it fitting to end this article by citing an old cliché used by Oppenheim as (for many of us) our DSP journey started with his textbooks.

1 + 1 = 3 [30]. Be creative.

Authors

Clive Cheong Took (clive.cheongtook@rhul.ac.uk) received his doctorate from Cardiff University. He is a senior lecturer at Royal Holloway, University of London, Egham, Surrey, TW20 0EX, U.K. His interest in teaching started in 2012 when he joined the University of Surrey. Since then, he became a fellow of the Higher Education Academy and was recently awarded a college excellence teaching commendation in recognition of his innovative project-led learning approach by the principal at Royal Holloway, University of London. He is an active member of the IEEE Signal Processing Society and co-organized the IEEE International Workshop on Machine Learning for Signal Processing in 2013 in Southampton and ICASSP 2019 in Brighton. He is a Senior Member of IEEE.

Stephen R. Alty (steve.alty@rhul.ac.uk) received his Ph.D. degree in signal processing from Liverpool John Moores University in 1998. In 2002, he became a lecturer in the Division of Engineering, King's College London. In 2011, he moved to the School of Engineering, London South Bank University. Since 2017, he has been a senior lecturer and director of undergraduate education at Royal Holloway, University of London, Egham, Surrey, TW20 0EX, U.K. His research interests include signal processing applications in biomedical engineering and pattern recognition. He is a member of the IEEE Signal Processing Society and IEEE Engineering in Medicine and Biology Society. He is a Senior Member of IEEE.

Anush Yardim (anush.yardim@rhul.ac.uk) received her Ph.D. degree from the University of Westminster, London. She is a senior lecturer at Royal Holloway, University of London, Egham, Surrey, TW20 0EX, U.K. After many years teaching at the undergraduate, postgraduate, and industrial professional development levels in the area of signal processing and communications, she joined Royal Holloway, University of London to help form a new electronic engineering department. Her research interests include innovation in engineering education and model-based digital signal processing learning. She is a chartered engineer, a fellow of the Institution of Engineering and Technology, a Member of IEEE, a senior fellow of the Higher Education Academy, and a fellow of the European Entrepreneurial Education Programme.

David M. Howard (david.howard@rhul.ac.uk) received his Ph.D. degree in human communication working on cochlear implants at University College London. He is a professor and founding head of the Department of Electronic Engineering at Royal Holloway, University of London, Egham, Surrey, TW20 0EX, U.K. His research interests include human voice production and perception in speech and singing, stimulated in part by being an organist and choir director. His Vocal Tract Organ uses 3D printed human vocal tracts as pipes that recently included that of the 3,000-year-old mummy, Nesyamun. He is a fellow of the Royal Academy of Engineering and the Institution of Engineering and Technology and a member of the Education and Skills Committees in both organizations.

References

[1] D. M. Howard, S. Kuenzel, and W. Liu, "From creativity and ingenuity through technology and invention to product and market: A new paradigm for engineering education." Presented at the Conf. New Approaches Eng. Higher Educ., May 2017.

[2] A. Karpathy. "Building the software 2.0 stack," presented at the Conf. Spark + AI Summit, 2018. [Online]. Available: https://karpathy.ai/ (accessed Jan. 19, 2021).

[3] J. P. Guilford, *The Nature of Human Intelligence*. New York: McGraw-Hill, 1967.

[4] "One kind of education does not fit all: Time for more variety, experimentation and creativity." The Economist. Oct. 2017. https://www.economist.com/special-report/2017/10/02/one-kind-of-education-does-not-fit-all

[5] S. R. Daly, E. A. Mosyjowski, and C. M. Seifert, "Teaching creativity in engineering courses," *J. Eng. Educ.*, vol. 103, no. 3, pp. 417–449, York, U.K., 2014. doi: 10.1002/jee.20048.

[6] N. Jackson, "Creativity in engineering education," The Higher Education Academy, York, U.K., 2014. [Online]. Available: http://www.creativeacademic.uk/uploads/1/3/5/4/13542890/creativity_in_engineering.pdf

[7] E. Hong, S. A. Hartzell, and M. Greene, "Fostering creativity in the classroom: Effects of teachers' epistemological beliefs, motivation, and goal orientation," *J. Crea. Behav.*, vol. 43, no. 3, 2009. doi: 10.1002/j.2162-6057.2009.tb01314.x.

[8] N. P. Kirillov, E. G. Leontyeva, and Y. A. Moiseenko, "Creativity in engineering education," *Procedia Soc. Behav. Sci.*, vol. 166, pp. 360–363, Jan. 2015. doi: 10.1016/j.sbspro.2014.12.537.

[9] Y. Tekmen-Araci and L. Mann, "Instructor approaches to creativity in engineering design education," *J. Mech. Eng. Sci.*, vol. 233, no. 2, pp. 395–402, 2019. doi: 10.1177/0954406218758795.

[10] DSP Creations. http://dspcreations.com/ (accessed Dec. 10, 2020).

[11] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing, New Int.'l Edition*, 3rd ed. London, U.K.: Pearson Education Ltd., 2014.

[12] H. W. Sun, A. Yardim, and G. D. Cain, "A modified filter bank for tracking multiple sinusoidal signals," *Proc. EUSIPCO 96*, vol. 1, pp. 691–694. Trieste, Italy, Sept. 10–13, 1996.

[13] G. D. Cain, H. W. Sun, and A. Yardim, "Two new cascade resonator-in-a-loop filter configurations for tracking multiple sinusoids," *Electron. Lett.*, vol. 30, no. 8, pp. 630–631, Apr. 1994. doi: 10.1049/el:19940436.

[14] Audacity. ver. 2.4.1. http://audacityteam.org/ (accessed June 15, 2020).

[15] D. Abercrombie, *English Phonetic Texts*. London: Faber and Faber, 1964.

[16] P. Boersma and D. Weenick. Praat: Doing phonetics by computer. 2018. http://www.praat.org (accessed June 6, 2020).

[17] H. Blair and B. May, "Sparky's Magic Piano and other stories," Capitol Records, 1947.

[18] M. Aylett and C. Pidcock, "The CereProc characterful speech synthesiser SDK," in *Proc. 7th IVA Intell. Virtual Agents*, Paris, 2007, pp. 413–414. doi: 10.1007/978-3-540-74997-4_65.

[19] M. S. Puckette, "Pure data," in *Proceedings of International Computer Music Conference*. San Francisco: International Computer Music Association, 1996, pp. 269–272.

[20] R. E. Remez, P. E. Rubin, D. B. Pisoni, and T. D. Carell, "Speech perception without traditional speech cues," *Science*, vol. 212, no. 4497, pp. 947–949, 1981. doi: 10.1126/science.7233191.

[21] D. Pressnitzer, J. Graves, C. Chambers, V. de Gardelle, and P. Egré, "Auditory perception: Laurel and Yanny together at last," *Curr. Biol.*, vol. 28, no. 13, pp. R739–R741, 2018. doi: 10.1016/j.cub.2018.06.002.

[22] S. T. Azeemi and S. M. Raza, "A critical analysis of chromotherapy and its scientific evolution," *J. Evid.-based Complem. Altern. Med.*, vol. 2, no. 4, pp. 481–488, 2005. doi: 10.1093/ecam/neh137.

[23] Panopto. https://www.panopto.com/ (accessed Nov. 29, 2020).

[24] Office for Students. https://www.officeforstudents.org.uk/ (accessed Nov. 29, 2020).

[25] S. Kanna, W. Von Rosenberg, V. Goverdovsky, A. G. Constantinides, and D. P. Mandic, "Bringing wearable sensors into the classroom: A participatory approach," *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 110–130, 2018. doi: 10.1109/MSP.2018.2806418.

[26] E. Richer and A. Nehorai, "Enriching the undergraduate program with research projects," *IEEE Signal Process. Mag.*, vol. 33, no. 6, pp. 123–127, 2016. doi: 10.1109/MSP.2016.2601652.

[27] H. Godrich, A. Nehorai, A. Tajer, M. S. Greco, and C. Zhang, "Special article series on signal processing education via hands-on and design projects," vol. 34, no. 1, pp. 13–15, 2017. doi: 10.1109/MSP.2016.2620199.

[28] J. Dul and C. Ceylan, "Work environments for employee creativity," *Ergonomics*, vol. 54, no. 1, pp. 12–20, 2011. doi: 10.1080/00140139.2010.542833.

[29] Y. Hadzigeirgiou, P. Fokialis, and M. Kabouropoulou, "Thinking about creativity in Science education," *Crea. Educ.*, vol. 3, no. 5, pp. 603–611, 2012. doi: 10.4236/ce.2012.35089.

[30] A. V. Oppenheim, "One plus one could equal three," *IEEE Signal Process. Mag.*, vol. 23, no. 6, pp. 10–12, 2006. doi: 10.1109/SP-M.2006.248707.

SP

Yolanda Campos-Roca

Multidisciplinary Project-Based Learning

Improving student motivation for learning signal processing

©SHUTTERSTOCK.COM/LUCKYSTEP

Due to its open-ended nature, project-based learning (PBL) is a suitable methodology to achieve student motivation and enrich creativity skills. This article describes the design of PBL activities in two undergraduate courses on digital signal processing techniques. The academic context is a bachelor's degree in telecommunication engineering at the Universidad de Extremadura in Spain. Full specifications for five project proposals that other instructors could use are explained, including an in-depth analysis of one of them in relation to the learning outcomes. The results of a survey designed to gain insight into students' perceptions of the PBL process are provided. All of the participants agree or strongly agree that PBL improves the long-term retention of knowledge and provides learning with a more practical orientation toward real goals than conventional learning. Also, 91% of them consider that PBL motivates them more than traditional methodologies. The instructor's conclusions after the project development and assessment are that PBL considerably helped the students to think more creatively and increased their motivation in comparison to other activities.

Introduction

In the current context of high pressure for innovation, the need for engineers with creativity skills is often emphasized [1], [2]. PBL is an attractive methodology to improve these skills [3]. Whereas many teaching–learning activities in engineering education focus on constrained problems leading to a unique right answer, open-ended PBL allows for the enrichment of divergent thinking skills. Also, student engagement increases when the projects refer to real-world applications.

Engineers must be able to work across disciplinary boundaries. Addressing PBL from a multidisciplinary perspective allows students to establish links among different areas of knowledge. Multidisciplinary projects also promote out-of-the-classroom discussions, giving the students the possibility to create connections with experts in other areas [4].

Signal processing has been defined by the IEEE Signal Processing Society as the "science behind our digital life." Therefore, signal processing courses are a suitable place to

Digital Object Identifier 10.1109/MSP.2021.3053538
Date of current version: 28 April 2021

 1053-5888/21©2021IEEE

incorporate a multidisciplinary perspective. While PBL methodology is not new in engineering education, there is much room in the literature for contributions about the design and implementation of challenging real-world, multidisciplinary projects in signal processing courses.

This article reports on the design of activities for two undergraduate signal processing courses at the Universidad de Extremadura. Detailed specifications for five multidisciplinary project proposals are provided, including an in-depth analysis of one of them in relation to the learning outcomes. Also, an assessment of the application of PBL in these two courses is presented, including students' perceptions obtained through a survey and the instructor's reflections.

Academic context

The context of this educational experience is a bachelor's degree in telecommunication engineering, accredited with the European Accredited Engineer label, EUR-ACE. This study program includes two compulsory courses on digital signal processing. The first course is Digital Processing of Audio and Video Signals (DPAVS) and provides six European Credit Transfer and Accumulation System (ECTS) credits. It is placed in the fifth semester and covers some of the fundamentals of digital signal processing: analog-to-digital conversion, discrete Fourier transform, the implementation of infinite-impulse response and finite-impulse response filters, and multirate signal processing.

Due to its open-ended nature, project-based learning (PBL) is a suitable methodology to achieve student motivation and enrich creativity skills.

The second course, named Digital Audio (DA), also provides six ECTS credits. It is placed in the seventh semester and focuses specifically on speech and audio signal processing. The contents in this case are: the fundamentals of speech and audio signal processing, analysis techniques, biomedical applications of acoustic signals, speech and audio compression, and a concise overview of multiple application areas of speech and audio signal processing.

Both courses are run using a combination of the following methodologies: expository lectures, paper and pencil exercises, guided practical sessions, and PBL activities. Guided practical sessions take place in a computer lab and have been designed to provide hands-on training on MATLAB programming to solve digital signal processing problems.

PBL-based methodology

A student-centered PBL methodology has been designed and applied based on the following aspects:

- open-ended project assignments
- real-world connection
- multidisciplinary context
- different project assignments that require overlapping knowledge
- time management through the use of Gantt charts
- multifaceted assessment through rubrics, in which creativity is included as a criterium.

Open-ended project assignments are designed to enrich divergent thinking skills. Only some specifications are given, while the rest of the project is open, encouraging the students to reach their own solutions. This is a key aspect in promoting their creativity and innovation capacity.

Students are organized into teams of two people, and several teams are grouped into a cluster. Each team develops a project. Each project assignment contains particular activities for clusters of teams (named "variations" in the assignments). The grouping of teams into clusters is different for each activity within the same project assignment. The proposal of similar (but not identical) projects to different teams allows for the establishment of two different collaboration levels: strong collaboration between the two students within each team and collaboration among different teams.

A factor that favors creativity is to propose a long-term project with a minimum time-frame of four weeks from the project proposal to the delivery of the written report and at least three additional days between the delivery of the report and the oral presentation. The generation of new ideas requires creative moments, which are often spontaneous and difficult to fit into a very tight time frame. Long-term projects also benefit students in terms of a better knowledge retention.

To be comfortable generating new ideas, a student needs "creative confidence" [5]. The facilitator's attitude is a major factor in developing this confidence by creating an environment where new ideas are never negatively criticized, especially in classroom discussions and oral presentations.

The design of project proposals

In the current context of a strong emphasis placed on educational technologies, effort should be invested by the instructors on deciding the challenges that students face. One of the key elements for a successful PBL methodology is a careful design of the project proposals in terms of the learning outcomes. Some general aspects about the design of the five project assignments presented are covered in the "General Aspects of Project Design" section, whereas the "Project Example: An Acoustic Rainfall and Hail Sensor" section gives an in-depth analysis of one of them.

The general aspects of project design

Project specifications are provided in "Project Topics" and the "Project Example: An Acoustic Rainfall and Hail Sensor" section. These specifications could be directly used by any instructor to run a PBL activity. The topics (and the course in which they are proposed) are the following: an acoustic rain detector (DA), audio steganography (DPAVS), spectral analysis of bird chirps (DPAVS), helium speech and formant synthesis (DA), and audio classification for hearing-aid applications (DA). See "Project Topics" for further details on these topics.

Project Topics

For extension reasons, this sidebar focuses on four project proposals. Other project-based learning assignments were designed and tested, for example, about the spectral analysis of marine sounds, audio fingerprinting with spectrograms, and digital music synthesis with low computational requirements.

The development of the projects requires a preliminary task related to time management and a final task consisting of preparing written reports and oral presentations, which are common to all of the proposals, as explained in the "Project Example: An Acoustic Rainfall and Hail Sensor" section. To avoid repetition, they are not included here.

Audio steganography

Steganography studies the techniques that allow for the concealment of messages so that their existence is not perceived. In contrast to cryptography, which focuses on making information understandable to any unauthorized person, steganography tries to hide the existence of the message. Specifically, a message of this type is concealed in another message that contains it (called a carrier), whose knowledge can be public and does not raise suspicions. Steganography can be used for legal and illegal purposes.

When the message is hidden in an audio file, the process is called audio steganography [S1]. An example technique of sound steganography consists of hiding an image in a spectrogram. This project focuses on this technique. The different activities to be performed are explained next.

Search for applications

The students should search for information and make a summary of the findings with an extension of 150–200 words. The topics depend on the cluster of teams. They include: for variation 1, a legal use of digital steganography, including source referencing; for variation 2, an illegal use of digital steganography, including a citation of an actual piece of news (not older than three years).

Search for an audio file

This activity consists of getting a suitable audio file. It must contain a hidden image that is visible through the calculation of a spectrogram (i.e., the audio file is the result of applying audio steganography). Each team can obtain such a file by directly performing an Internet search (files with Creative Commons license) or search for freely available code to create it. Proper source citing is mandatory. Each team will use different audio files. The coordination will be done in the classroom.

Spectral analysis on simple test signals

The student teams will generate spectrograms from audio signals hiding an image. However, to make the interpretation of these spectrograms easier, it is proposed that students perform a test with simple signals as a preliminary step. Using MATLAB, students should investigate the influence of segment length and window type on the results of the analysis. Time and frequency resolution will be taken into account.

The test signal will be created as the concatenation of two segments of cosenoidal signals, with amplitudes A_1 and A_2 and frequencies f_1 and f_2. The sampling frequency will be 44,100 Hz, and the length of each segment is N_{sample} samples. The teams should modify the frame length used in the short-term spectral analysis (keeping the window type constant) and also check the influence of the window type (keeping the frame length constant). The specifications for the first cluster of teams (variation 1) are: $N_{\text{sample}} = 12{,}000$; $A_1 = A_2 = 1$; $f_1 = 2{,}003\,\text{Hz}$, $f_2 = 2{,}503\,\text{Hz}$; rectangular and Blackman windows will be compared; to study the influence of segment length, rectangular windows will be used. For the second cluster of teams (variation 2), the specifications are: $N_{\text{sample}} = 9{,}000$; $A_1 = A_2 = 2$; $f_1 = 1{,}600\text{ Hz}$, $f_2 = 2{,}600\,\text{Hz}$; rectangular and Hamming windows will be compared; to study the influence of segment length, Hamming windows will be used.

Spectral analysis with real audio steganography signals

Once the main aspects of spectrogram analysis are understood through the use of simple signals, the students should perform a spectral analysis of the audio file containing the hidden image obtained before. In a similar way to that proposed in the aforementioned section, the students should investigate the influence of segment length and window type on the results.

Spectral analysis of bird chirps

A sinusoidal signal consisting of a frequency sweep over time is called chirp (https://es.mathworks.com/help/signal/ref/chirp.html). The frequency can increase or decrease, and the variation can be nonlinear. Also, the amplitude envelope does not have to be constant. In general, a signal is considered to be chirp-like if its amplitude envelope shows a slow evolution compared to its frequency variation.

Nature offers numerous examples of chirp-like sounds. For example, some animals (birds, bats, frogs, dolphins, whales, and so on) emit this type of sound, which is of great interest in the field of bioacoustics. By identifying specific acoustic features, it is possible to perform automatic species classification. In this project, spectral analyses will be performed on chirp-like signals recorded from birds.

Search for information about bird species emitting chirps

Extremadura is a region in Spain especially known for the conservation of some of the most threatened bird species in Europe. The different clusters will search for information about the following topics and make a summary of the findings with an extension of 150–200 words: bird species that are present in this region and emit chirp-like sounds (variation 1) and research studies where spectrograms of bird chirps have been used (variation 2).

Spectral analysis on simple signals

First of all, the main concepts involved will be worked out through the use of simple synthetic signals. A synthetic chirp signal will be generated with the following specifications: for variation 1, a linear frequency variation from 150 to 700 Hz in 2 s, at a sampling rate of 2,000 Hz; for variation 2, an exponential frequency variation from 200 to 800 Hz, in 1 s, at a sampling rate of 4,000 Hz.

A short-term spectral analysis will be performed on the synthetic signal to create a spectrogram. The impact of frame length and windowing type on the results should be analyzed. To do this, the window type will be fixed while the frame length is varied, and afterward, the frame length will be fixed, and the window type will be varied. The fixed window type will be rectangular (variation 1), Bartlett (variation 2), or Hamming (variation 3). The comparison will consider rectangular versus Blackman windows (variation 1), rectangular versus Bartlett windows (variation 2), and Hamming versus Blackman windows (variation 3).

Search for sound data files

To perform the next task, the students need a sound file that corresponds to a chirp-like sound of a bird. The teams will search for sound files available (with Creative Commons license) on the Internet, or they will make the recording themselves. An example of an online database is https://www.xeno-canto.org/, but it is not mandatory to use it. It is recommended that the sound produces an easy to interpret spectrogram (i.e., the frequency sweep can be seen as clearly as possible). This depends not only on the sound but also on the parameters chosen for the spectral analysis so this section should be considered together with the next one.

Spectral analysis on bird sounds

A short-term spectral analysis will be performed on the bird signal, and the spectrograms will be interpreted. The impact of the frame length and windowing type on the results will be analyzed. The final report will include a discussion about all of the results. Window types should be the following: a Tukey window with different values of its cosine factor parameter (variation 1), a Kaiser window with different values of its shape factor (variation 2), and a Chebyshev window with different values of its sidelobe magnitude factor (variation 3).

Helium speech and formant synthesis

This project deals with the analysis of helium speech as a context for learning about the source–filter model of speech production and formant synthesis of speech. The helium effect is easy to understand by considering the source–filter model of speech production. Inhaling helium changes the formants of speech (and therefore the spectrum envelope), but it does not change the pitch.

Although the vocal tract shape can be complex, its effect can be modeled by one or more acoustic tubes. The simplest model is based on a single cylindrical tube with a uniform cross section that is closed at the location of the vocal folds and open at the mouth. A more accurate model can be obtained by connecting two or more uniform tubes with different cross-sectional areas [S2]. The resonance frequencies depend on each particular vocal tract geometry. For a tube with length L, the lowest resonant frequency takes the value $F_1 = c/4L$, where c is the speed of sound in the tube. Due to this dependence on the speed of sound, formants are considerably affected by helium. Helium is lighter than air, and sound waves travel more quickly through it.

Because it sounds very funny, helium voice often appears up in movies and television shows. "The Helium Insufficiency" episode of the popular television series *The Big Bang Theory* is a well-known example. Deep-sea divers also breathe in an atmosphere based on a helium-rich gas mixture, with a consequent impact on their voices. We will use a fragment of helium speech and a fragment of normal speech in air to analyze the helium effect. After the analysis tasks, some synthesis tasks will be addressed. Formant synthesizers are based on the source–filter theory of speech production. The transfer function of the vocal tract is modeled by a connection of simple resonators simulating formant frequencies and bandwidths. The proposed synthesis tasks refer to this technique.

Search for information about helium speech unscramblers

The teams will search for information about helium speech unscrambler technology available for communication with sea divers. They should describe a commercial product from one manufacturer and give a technical description in 150–200 words. Different clusters of teams will focus on manufacturers from different world regions: Europe and Africa (variation 1), North and South America (variation 2), and Asia and Australia (variation 3).

Analysis of speech in helium versus speech in air

The teams need sound samples corresponding to the same speaker before and after inhaling helium from a balloon. It is relatively easy to find videos on the Internet showing the experiment (for example, on YouTube). The students should not perform the experiment themselves but search for available sound files or videos. Using the filters available for advanced search, only videos freely offered under a Creative Commons license will be considered.

(*continued*)

Project Topics

Vowel segments will be considered. The specifications for the different clusters are the following: for variation 1, vowels/æ/(b**a**t) or/ Δ /(b**u**t); for variation 2, vowels/ ∂ /(**a**head) or/ э :/(h**e**r); for variation 3, vowels/i:/(b**ea**t) or/ Ι /(b**i**t); for variation 4, vowels/u/(b**oo**t) or/ ℧ /(b**oo**k).

Formant extraction

The formant frequencies of different vowels obtained from ordinary speech and helium speech will be compared. The results will be compared with the theoretical predictions made with acoustic tube models. Only the concatenation of simple lossless tubes with uniform cross sections will be considered. To extract the values of the formant frequencies from a speech segment, it is suggested that students calculate the roots of the linear prediction polynomial.

Formant-based synthesis of vowels

The formant synthesis of speech is based on a set of rules that control a highly simplified source–filter model, where the vocal tract characteristics are modeled by the interconnection of second-order all-pole filter sections, one per formant. These second-order resonators can be combined in parallel or cascade configuration. In this section, we focus on vowel sound synthesis, for which the cascade approach is more suitable. In a cascade configuration, there is no need to specify gains for each formant resonator, and two parameters only are necessary for each single section: the resonance center frequency and the bandwidth. Resonator coefficients can be computed from these two parameters and the sampling rate.

Cascade formant synthesis of stationary vowels

The next step will be to implement a MATLAB function that generates a stationary vowel sound using formant synthesis. The function will allow students to model speech in air or speech in helium. The necessary specifications are the fundamental frequency, duration, sample rate, and concrete vowel. The input to the first resonator will be the glottal source, modeled as a pulse train (Rosenberg model [S3]). Thereafter, the input to a resonator is the output of the previous one. Also, a radiation model will be added after the resonators. After implementing the function, the vowels specified for each cluster of teams will be synthesized in air and helium. Also, the teams will investigate the influence of the number of formants considered in the final quality.

Formant synthesis of vowel sequences

In the previous task, stationary vowels were synthesized. In this section, the teams will go a step further by modeling a time-varying vocal tract so that sequences of vowels can be generated. To get a natural transition, the formants of one vowel must gradually change in small steps to the ones of the subsequent vowel so that an abrupt discontinuity is avoided. Each team will choose the parameter values for the algorithm so that the transition is as natural as possible. The algorithm will be tested by generating several vowel sequences in which the formant frequencies and bandwidths gradually change over time.

eSpeak text-to-speech converter

eSpeak is an open source software speech synthesizer based on a formant synthesis method. After downloading the software (http://espeak.sourceforge.net/), the teams will execute it to obtain speech in Spanish and test the intelligibility of the synthetic speech by using the following subjective tests: a simple version of the Diagnostic Rhyme Test based on a small number of word pairs (variation 1) or a simple version of the Modified Rhyme Test based on a small number of word pairs (variation 2).

Audio classification in the context of hearing aid applications

Different acoustic environments require different hearing aid settings to achieve the best sound quality. Therefore, hearing aids usually include a module to classify sounds

The topics were selected to engage the students. Some of the projects in the DA course involve the use of machine learning techniques. The combination of signal processing and machine learning has a lot of interest in today's marketplace since it allows for addressing many different multidisciplinary projects.

These PBL activities are aimed at contributing to the achievement of different transverse and specific learning outcomes. The transverse ones are common for all of the project proposals, although the level of acquisition pursued is different depending on the course. They are the following: to synthesize and extract the necessary information to solve a problem, to work in an active and autonomous way, to develop creative solutions, to develop a lifelong learning mindset, to work collaboratively, to learn technical vocabulary in English (since the students are Spanish native speakers), to communicate results through the preparation of technical reports and oral defenses, and to improve time management skills.

The specific learning outcomes are: to know the applications of spectral analysis, to understand the impact of segment length and window type on a spectral analysis, to apply feature extraction techniques from audio signals, to understand rule-based speech synthesis, and to produce programs in MATLAB. See "Project Topics" for more details.

All of the aforementioned learning outcomes are related to program outcomes and match the institution's vision of developing "a model of its own identity with quality teaching" and sensed by the microphone. Then, the sound is processed using different algorithms depending on the result. Classification systems consider a set of features, carefully chosen to emphasize signal characteristics that allow it to perform a discrimination. The goal of this project is to develop a classification system of sound types based on wavelet features.

Search for information about commercial hearing aids

The student teams will search for information about sound classification used in hearing aids in the current market. Although it can be difficult to find out all of the technical details of a commercial product, it will always be possible to provide an explanation of functionality from the user's point of view. The technical description will be as deep as the availability of public information allows. The extension should be 150–200 words. The different clusters of teams will look for commercial products of manufacturers from different world regions: Europe and Africa (variation 1), North and South America (variation 2), and Asia and Australia (variation 3).

Preparation of the sound database

Sound samples will be necessary to check the performance of the system. There are multiple sound databases freely available; however, each team will create its own database. It is recommended that they take into account the daily life situations of people using hearing aids.

Feature extraction

A major step in the design of a classification system is the selection of an efficient set of features that are capable of discriminating the signals. Short-time Fourier transform is traditionally used in audio feature extraction for time–frequency decomposition. The main drawback is its resolution limitation, which is dependent on the window size: the shorter the analysis window, the better the time resolution, but the poorer the frequency resolution.

To overcome the limitation of a fixed analysis window, the wavelet transform uses short windows at high frequencies and long windows at low frequencies. The result is that different resolutions are obtained for different frequencies. For this reason, this type of analysis is known as multiresolution analysis [S4].

In this project, a decomposition of the audio signal based on the discrete wavelet transform will be used. Using the MATLAB wavelet toolbox (https://es.mathworks.com/products/wavelet.html), it is possible to divide the spectrum of the signals into sub-bands. Then the energy distribution in each sub-band will be calculated. These energies are the features on which the classification is based. The number of extracted features depends on the number of levels of decomposition performed in the wavelet transformation.

The MATLAB wavelet toolbox includes several families of wavelets that are commonly used. The specifications for the different clusters of teams are the following: Daubechies (variation 1), Symlet (variation 2), and Coiflet (variation 3). The different teams will investigate the impact of different decomposition levels and wavelet order on the classification performance. A discussion will be included in the report.

Classification

This classification task is proposed in the same way as in the "Project Example: An Acoustic Rainfall and Hail Sensor" section.

References

[S1] F. Djebbar, B. Ayad, K. Abed Meraim, and H. Hamam, "Comparative study of digital audio steganography techniques," *EURASIP J. Audio, Speech, Music Process.*, vol. 2012, no. 1, pp, 1–16, 2012. [Online]. Available: https://doi.org/10.1186/1687-4722-2012-25

[S2] C. D. Wentworth, "Helium speech: An application of standing waves," *Phys. Teacher*, vol. 49, no. 4, pp. 212–215, 2011.

[S3] A. E. Rosenberg, "Effect of glottal pulse shape on the quality of natural vowels," *J. Acoust. Soc. Amer.*, vol. 49, no. 2B, pp. 583–590, 1971.

[S4] E. Didiot, I. Illina, D. Fohr, and O. Mella, "A wavelet-based parametrization for speech/music discrimination," *Comput. Speech Lang.*, vol. 24, no. 2, pp. 341–357, 2010.

providing "a firm and resounding orientation toward research, innovation and its consequent transfer." They also match the University of Extremadura's mission statement, which includes the following aspects: "integral formation," "quality teaching and research," "key role in the development of the regional society," "pursuit of excellence," and "national and international projection."

PBL should help students establish connections to the real world. Therefore, all of the project proposals involve the processing of realistic signals. In the case of the lower course, the project assignments include a first activity based on ideal synthetic signals. The motivation behind this is that some signal processing concepts are easier to understand if they are first applied to an ideal signal. This is just a preliminary step in the projects that are always focused on an application with realistic signals.

Project example: An acoustic rainfall and hail sensor

This section presents an example of a project specification and discussion. The assignment was used in the DA course in the academic year 2019–2020, when the student survey reported in the "Students' Perceptions of PBL" section was carried out. The different aspects taken into account in the design are analyzed. The paragraphs included in the students' assignments are shown in italics.

The project deals with the development of algorithms for automatic rainfall and hail detection from acoustic recordings by using signal processing and machine learning. The assignment belongs to the area of pure PBL since there is not a theoretical lecture about machine learning techniques in the course, and, thus, the students learn through the project. Also, some of the acoustic features used were not previously explained in a theoretical lecture.

Apart from other goals usually present in PBL activities, this proposal was conceived with the aim of promoting sustainable and socially responsible engineering. It was taken into account that the United Nations' Sustainable Development Goals (SDGs) should be increasingly present within the universities' learning programs in an effort to reach the United Nations 2030 Agenda. The assignment presents this motivation:

Due to climate change, the risk of flooding is increasing. In this context, rainfall and hail sensors are important as they can provide an alarm system for areas highly prone to floods or landslides. Placing these sensors would allow to detect if rainfall or hail rates are exceeding threshold levels.

In the scientific literature, it is possible to find several contributions on acoustic rainfall sensing systems. These systems are based on recorded sound produced by impact of raindrops or hail particles and provide near real-time data. Most of current acoustic systems for rainfall detection have been proposed for sound recorded with hydrophones in the sea [9]. *Some examples of low-cost and portable acoustic sensors can also be found in terrestrial environments* [10].

Power distribution of rainfall and hail sound through the spectrum is not the same for different environments. For these reasons, the scenario must be taken into account. Also, the sound is different depending on the precipitation intensity.

In this project, a simple system for automatic detection of heavy rainfall and hail events by using acoustic recordings in a terrestrial environment will be designed and coded in MATLAB (see Figure 1). *Also, the students should provide some ideas about how to implement a complete detection system for a concrete application.*

Creativity is fostered on several occasions through the assignment. The first time the assignment text refers to creative solutions is the following:

The environment to consider depends on the cluster of teams, as follows: forest environment (variation 1), urban environment (variation 2). Each team can freely choose the specific characteristics of the forest or urban environment. This is an open-ended project, where many different systems are possible. Please feel free and use your creativity to develop your own system.

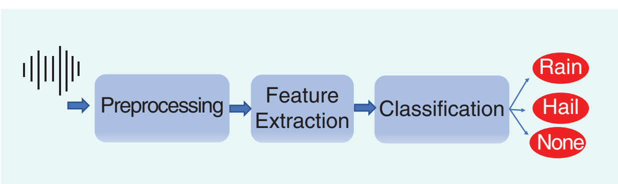

FIGURE 1. The sound classification system.

The development of the project includes the preparation and update of a project schedule. The requirement of a Gantt chart one week after proposing the project is proven to have improved students' time-management capabilities. It is proposed as follows:

Project schedule needs to be planned at the beginning. This involves creating a list of project tasks (and student assigned, for team work efficiency) and project timeline, with the aid of Gantt charts. The initial plans might suffer modifications so that the final Gantt chart might be different from the initial one. These deviations should be reflected in the final report.

Connection with the real world is emphasized by promoting the realization of realistic own recordings to create the sound database:

One task of this project is to create a sound data set. This dataset will include rainfall and hail sounds, completely different sounds and also other sounds that could be easily confused with rain or hail. The files can be, if possible, recorded by the students or can be downloaded from data sets (under Creative Commons license) from the Internet.

The different learning styles and capabilities that students have must be addressed by projects that allow extensions with varying levels of difficulty. For example, in the creation of the database, different solutions were allowed: own recordings, audio files downloaded from the Internet, or hybrid solutions. Due to rainfall scarcity (especially hail) in the Spanish region where the university is located, some teams used a hybrid approach (own recordings combined with downloaded files), whereas other teams directly chose the easiest way, i.e., to download files from the Internet. The latter is also useful for increasing the ability to search for and analyze information. Potential differences in the recording conditions of the sound samples and their potential impact on the results were analyzed only by some teams.

Many possible features could be used to perform class discrimination. Figure 2 plots a spectrogram displaying spectral differences between rain and hail sound segments as well as a third segment in which only traffic noise was present. The motivation to propose features was that the extraction process was based on simple algorithms that are easy to understand by the students. The suggested features were band energy ratio and the following spectral descriptors: centroid, spread, skewness, kurtosis, entropy, flatness, crest, flux, slope, decrease, and roll-off point. These descriptors were not specifically defined and explained in a theoretical lecture but learned through the project. The project is open ended, and different teams used different feature subsets as the assignment proposed:

Features are any measurable aspect of an audio signal that might be used for classification purposes. A major step in the design of a signal classification system is the selection of an efficient set of features that are capable *of discriminating the classes. In general, the magnitude of the spectrum increases with rainfall intensity. However, this increment is more noticeable in specific frequency bands. The band energy ratio (the ratio of the energy in specific frequency-bands to the total energy) can be used as a feature to discriminate sounds. Other spectral features than can be used for classification purposes can be tested (https://www.mathworks.com/help/audio/ug/spectral-descriptors.html). Each team can choose a suitable set of features that are able to discriminate the different classes.*

The preferred features chosen by most of the teams were spectral centroid and roll-off, followed by band energy ratio. Two features not mentioned in the assignment—but theoretically studied in the lectures—were short-term energy and zero-crossing rate, and these were included by more than one half of the teams.

Multidisciplinarity is introduced through the application of machine learning techniques. It is used as a tool to increase student motivation. This new topic could be difficult to work out by a single team; however, assignments for different teams require some overlapping knowledge. During project development, there will be discussions about machine learning involving different teams. This has proven to be very effective in the learning process of a new topic. Therefore, both learning outcomes (i.e., to work collaboratively and to synthesize and extract the necessary information to solve a problem) are addressed. Also, creativity is fostered again by the last two sentences in the following paragraph of the assignment:

Classification experiments will be conducted to check system performance. Pattern recognition techniques can be used to automatically classify objects into categories using a training dataset. There are multiple classification engines that could be tested. However, since this course is mainly about digital audio processing, it is suggested that the students invest more effort in the feature extraction process and focus on a concrete type of classifier. Discriminant analysis or support vector machines are recommended. For predictive accuracy estimation, a leave-one-out cross-validation method is suggested. Each team will decide and set-up its own classification experiments. Creative ideas will be highly appreciated.

Although variations were accepted, all of the teams decided to use discriminant analysis [7] and, concretely, the classify function in MATLAB. The probable reason was that the project introduced a completely new knowledge field, and students did not feel self-confident enough to introduce

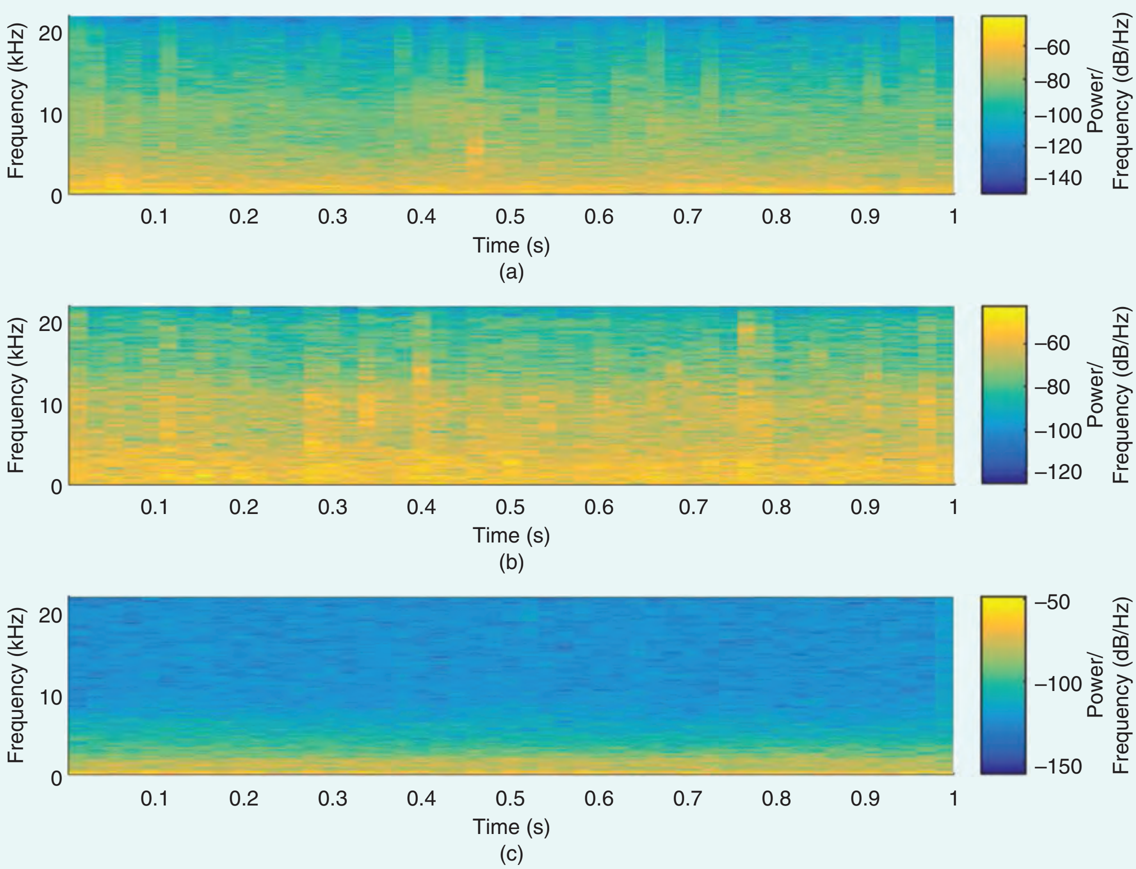

FIGURE 2. The spectrograms of different sounds [6]. (a) Rain, (b) hail, and (c) no rain or hail, just traffic.

variations. Since the project is proposed in an audio signal processing course, machine learning tasks are proposed in a complementary way, and it is difficult to go deep in this area. For the same reason, automatic feature selection was avoided, and the instructor simply recommended the use of a small feature subset selected by the students. Some creative solutions were provided, for example, by setting an additional binary discrimination task where rain and hail sounds constituted a single precipitation class.

Lab projects often have a poor connection with the real world because they lack a system perspective. For this reason, it is required that the students conceive the whole system. This task also promotes creativity because students can propose their own application scenarios. The text in the assignment is as follows:

Some ideas about how to implement the complete detection system for a concrete application should be included as future directions.

To become aware of new challenges that it is possible to face in a field provides motivation to create a lifelong learning mindset. With this goal in mind, the following task is included in the assignment:

Each team should propose a new topic in the area of environmental signal processing.

Concerning the development of communication skills, the following tasks were proposed:

After completion of the experiments, the following deliverables should be uploaded to the Moodle virtual classroom: a technical report prepared according to a given format (IEEE double-column template for conference papers) with a limit of 6 pages; slides for a 10 minute (5 minute/person) oral presentation; the source code and any data that is necessary to replicate the results.

Oral presentations in small groups (six students and the professor compose the audience) are a key aspect to building self-confidence in a low-risk environment.

The key role of grading rubrics as assessment tools

Grading rubrics are key tools to help instructors and students focus their attention on the most relevant aspects of an assignment. Assessing creativity is particularly difficult because it is a subjective concept. However, it is necessary to define some subcriteria so that students know that the risk of making mistakes by trying innovative ways will be recognized in the final grade. Creative thinking is assessed through the following subcriteria:

- propose ideas that might be uncommon or risky but that show some potential to become useful
- propose and test multiple solutions and critically evaluate their performance
- propose changes to the assignment specifications to enrich the project facing new challenges
- relate the challenge to other situations in different application contexts
- use original visually engaging elements in the oral presentation.

Two rubrics are used to evaluate the project assignments in both courses, one for the report and another one for the oral presentation. The first one includes the following criteria: results, depth of analysis, conclusion, bibliography, structure and organization, use of equations, figures and tables, format, language (since students are Spanish, writing in English is positively assessed), quality of writing (word choice, spelling, and grammar), teamwork, creativity, and innovation.

The rubric used for the oral presentation includes the following assessment criteria: subject knowledge, organization, quality of the slides, visually engaging elements, language (if it is written in English), delivery techniques (elocution, enthusiasm, posture gestures, and eye contact).

Even in the case of the higher-level course, where many students have good MATLAB programming skills, 36% of them would increase the number of guided practical sessions.

Students' perceptions of PBL

During the academic year 2019–2020, the students responded to a questionnaire designed to collect information on their perceptions about this PBL methodology. The survey was implemented online within the institutional Moodle virtual classroom environment. Participation in the survey was voluntary and anonymous. The target groups were 27 students in the DPAVS course (13 of which voluntarily participated in the survey) and 22 students in the DA course (with 11 voluntary participations). In the DA course, the project was the rainfall detection system described in the "Project Example: An Acoustic Rainfall and Hail Sensor" section.

The results of questions 1–3 (Figure 3–5) illustrate that the participants mostly think that PBL increases their motivation for the subject (91% agree or strongly agree), gives the learning experience a more practical orientation toward real objectives (100% agree or strongly agree), and improves the long-term retention of knowledge (100% agree or strongly agree). In the case of DA, these results confirm that the strategies followed in the design of the rainfall and hail detection project to achieve students' engagement had fruitful results.

In the case of question 3 (Figure 5), the results obtained from the fourth-year students are considered more reliable because they experienced the methodology one year before and thus had a better perspective of the long-term retention of knowledge.

According to the responses to question 4 (Figure 6) in both courses, 70% of the students would reduce (62%) or eliminate (8%) the number of expository lectures.

The answers to question 5 (Figure 7) reveal that guided lab sessions are considered necessary for most of the students, which means that students expect some initial guided practical training. Even in the case of the higher-level course, where many students have good MATLAB programming skills, 36% of them would increase the number of guided practical sessions.

The answers to question 6 (Figure 8) reinforced the idea that the PBL methodology is useful since 96% of the students

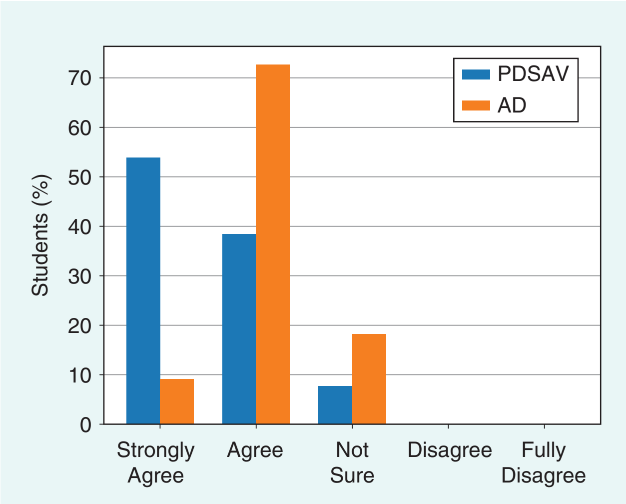

FIGURE 3. The distribution of responses to question 1: PBL increase students' motivation for the subject.

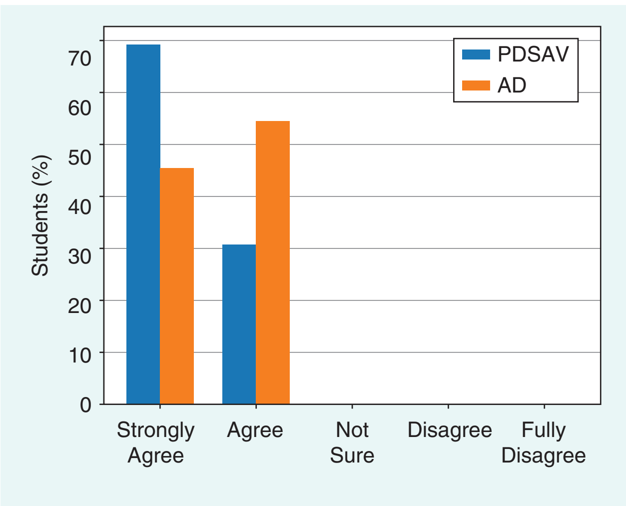

FIGURE 4. The distribution of responses to question 2: PBL gives the learning experience a more practical orientation toward real goals.

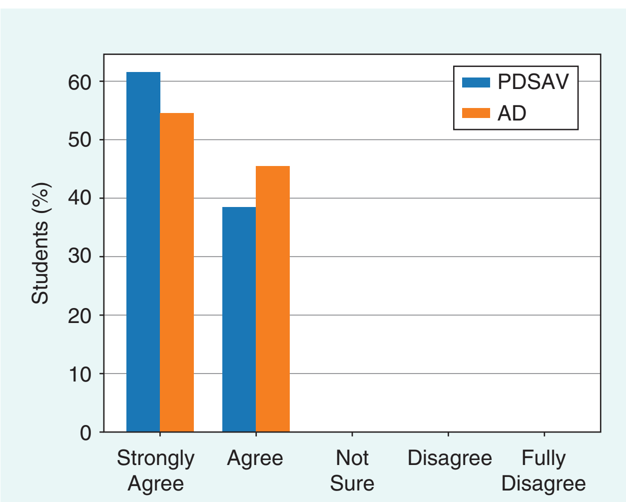

FIGURE 5. The distribution of responses to question 3: PBL improves the long-term retention of knowledge.

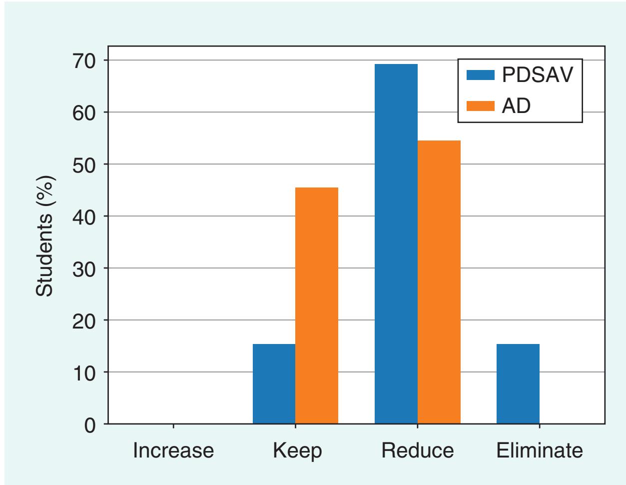

FIGURE 6. The distribution of responses to question 4: Assuming the total workload of this course does not change, I would … the number of expository lectures.

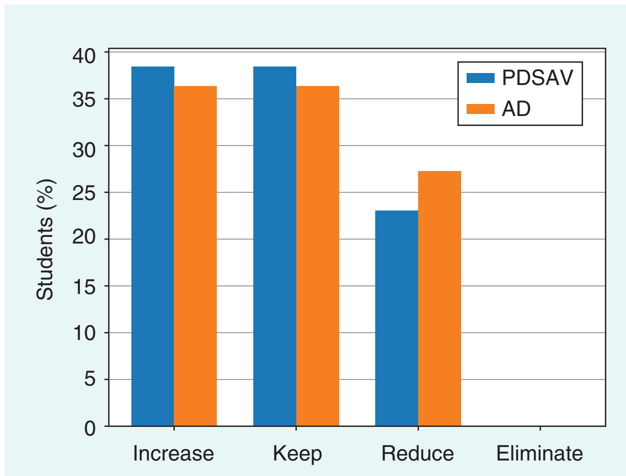

FIGURE 7. The distribution of responses to question 5: Assuming the total workload of this course does not change, I would … the number of guided practical sessions.

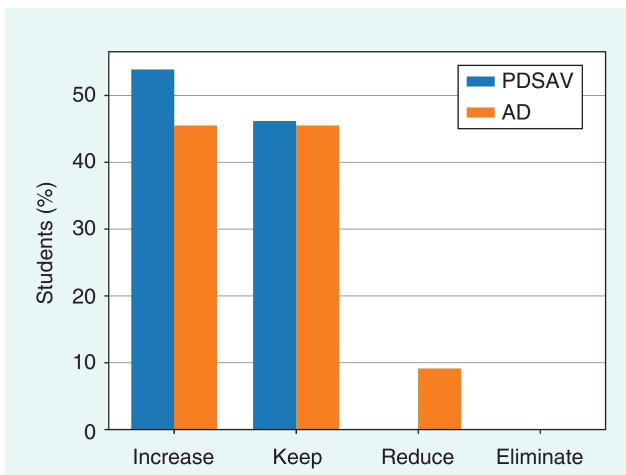

FIGURE 8. The distribution of responses to question 6: Assuming the total workload of this course does not change, I would … the current workload of PBL activities.

would keep or increase the workload of this type of activity. Only one student would reduce the workload, and none of them would eliminate the PBL activities.

Discussion

The considered courses are based on a combination of methodologies. A pure PBL scenario has not been tested. The author considers that some traditional instruction is still necessary at the undergraduate level to provide some background information and basic concepts. One of the main difficulties in applying pure PBL, especially in the fifth semester, is that students are not sufficiently used to this methodology. This is also the students' perception since only two out of the 24 surveyed students considered that expository lectures should be completely removed.

Guided lab sessions have a high formative value because the students actively participate. Since they try to finish the tasks within the session, they don't lose concentration. Many university courses rely entirely on this type of activity. As previously mentioned, due to the limited time schedule and closed-ended nature of the exercises, there is little place for divergent thinking.

PBL has proven to be very effective in meeting the diverse learning needs of all of the students. Students have different levels of motivation, learning styles, and responses to instructional practices. Student-centered methodologies are very powerful in facing the scenario of student diversity. PBL has also demonstrated its effectiveness concerning the enrichment of divergent thinking skills. Similar challenges have led to a wide variety of solutions.

Since the DA course uses knowledge from the DPAVS course, the author has been able to positively assess the retention of knowledge from the instructor's point of view. The author's perceptions are in agreement with the students' responses. The concepts and techniques related to short-term spectral analysis have been assimilated by the students, and they remember them the next year.

Apart from the aforementioned benefits, the experiences have allowed for the identification of a positive collateral result. When students perform traditional learning activities or even innovative activities, but in the disciplinary context of their engineering field, they do not speak very much about it outside of the academic context. In contrast, multidisciplinary projects are more susceptible to dissemination. Currently, there is a shortage of engineering vocations. Multidisciplinary PBL helps to create a learning environment that transcends the classroom. This should contribute to explaining the role of engineers in society and consequently engage more young talent in engineering studies [8].

Conclusions

PBL methodology has been successfully applied to teach signal processing in two undergraduate courses at the Universidad de Extremadura in Spain. Survey results have demonstrated that students have a positive view of these activities.

Due to the benefits of cthis methodology, initiatives aiming at sharing PBL experiences and material should produce a positive impact on the digital signal processing teaching–learning community. Future work should focus on the design and sharing of new PBL resources with a social dimension that helps to explain the important role of engineers in society within the framework of the United Nations SDGs.

Acknowledgments

The author would like to thank all of the students who took part in the survey. This work has been supported by project GR18055 (Junta de Extremadura/European Regional Development Funds, European Union).

Author

Yolanda Campos-Roca (ycampos@unex.es) received her M.S. and Ph.D. degrees in telecommunication engineering from the University of Vigo, Spain, in 1994 and 2000, respectively. In 2000, she joined the Department of Computer and Communication Technologies, Universidad de Extremadura, Cáceres, 10003, Spain, as an assistant professor, and in 2002, she became an associate professor. Her research interests focus on two different fields: speech signal processing for biomedical applications and circuit design in microwave, millimeter-wave, and submillimeter-wave bands. Her teaching activities are linked to her research and also involve two areas: digital signal processing and high-frequency devices and circuits.

References

[1] A. Valentine, I. Belski, M. Hamilton, and S. Adams, "Creativity in electrical engineering degree programs: Where is the content?" *IEEE Trans. Edu.*, vol. 62, no. 4, pp. 288–296, 2019. doi: 10.1109/TE.2019.2912834.

[2] T. Haertel, C. Terkowsky, and M. Radtke, "Creative students need creative teachers: Fostering the creativity of university teachers: A blind spot in higher engineering education?" in *Proc. Int. Conf. Int. Collab. Learn. (ICL)*, 2015, pp. 137–140.

[3] A. Isabekov and G. Sadyrova, "Project-based learning to develop creative abilities in students," in *Vocational Teacher Education in Central Asia*. Cham: Springer-Verlag, 2018, pp. 43–49.

[4] T. Schäck, M. Muma, and A. M. Zoubir, "Signal processing projects at Technische Universität Darmstadt: How to engage undergraduate students in signal processing practice," *IEEE Signal Process. Mag.*, vol. 34, no. 1, pp. 16–30, 2017. doi: 10.1109/MSP.2016.2619218.

[5] T. Kelley and D. Kelley, *Creative Confidence: Unleashing the Creative Potential Within Us All*. Currency, 2013.

[6] F. Font, G. Roma, and X. Serra, "Freesound technical demo," in *Proc. 21st ACM Int. Conf. Multimedia*, 2013, pp. 411–412. doi: 10.1145/2502081.2502245.

[7] T. Li, S. Zhu, and M. Ogihara, "Using discriminant analysis for multi-class classification: An experimental investigation," *Knowledge Inform. Syst.*, vol. 10, no. 4, pp. 453–472, 2006. doi: 10.1007/s10115-006-0013-y.

[8] D. López Álvarez, "Using service learning for improving student attraction and engagement in stem studies," in *Proc. 45th SEFI Conf.,* Azores, Portugal, 2017, pp. 1053–1060.

[9] E. Amitai and J. A. Nystuen, "Underwater acoustic measurements of rainfall", in *Precipitation: Advances in Measurement, Estimation and Prediction*. Berlin: Springer-Verlag, 2008, pp. 343–363.

[10] C. Bedoya, C. Isaza, J. M. Daza, and J. D. López, "Automatic identification of rainfall in acoustic recordings," *Ecological Indicators*, vol. 75, pp. 95–100, Apr. 2017.

SP

Meinard Müller, Brian McFee, and
Katherine M. Kinnaird

Interactive Learning of Signal Processing Through Music

Making Fourier analysis concrete for students

In this artaicle, we illustrate how music may serve as a vehicle to support education in signal processing. Using Fourier analysis as a concrete example, we demonstrate how the music domain provides motivating and tangible applications that make learning signal processing an interactive pursuit. Furthermore, we indicate how software tools, originally developed for music analysis, provide students multiple entry points to delve deeper into classical signal processing techniques while bridging the gap between education and cutting-edge research.

Introduction

Music is a ubiquitous and vital part of our lives. Thanks to the proliferation of digital music services such as Spotify, Pandora, and iTunes, we can enjoy music anytime and anywhere, interacting with it in a variety of ways, both as listeners and active participants. Aside from human speech, music may be the most familiar form of structured audio to most people. Conversely, as a scientific discipline, signal processing can be obtuse and unfamiliar to newcomers. The conceptual and practical understanding of signal processing requires a rather sophisticated knowledge of advanced mathematics, which can make the subject intimidating—even at the introductory level.

In this article, we show how music may serve as a vehicle to make learning signal processing an interactive pursuit, whether through concrete examples, hands-on exploration, or experimentation. The inclusion of music bridges the gap between the humanities and more typical signal processing communities such as mathematics, computer science, and engineering. This is the reason why one can find music as an integral part of books on multimedia and audio signal processing [1], [2]. Throughout this article, we show how music yields an intuitive entry point to support education on various levels. This leads to a learning approach that Guzdial [3] calls a "contextualized educational experience" for signal processing.

This article presents a scaffold for incorporating interactive music-based examples and music technology into an existing signal processing course. The proposed pipeline moves students through Bloom's taxonomy (Figure 1), helping them transition

Digital Object Identifier 10.1109/MSP.2021.3052181
Date of current version: 28 April 2021

from passive learners to engaged researchers and practitioners [4], [5]. This scaffolding resembles "legitimate peripheral participation" [6], aligning classroom learning more closely to the apprentice model of how signal processing is actually researched and practiced. Presented as a series of small but meaningful steps, this scaffold adds music to a signal processing course, transitioning it from a standard lecture into one that is interactive and project based. Each step provides a more context-rich signal processing course than the previous step. The examples in the presented scaffold are informed by the field of music information retrieval (MIR), which has interests in extracting the semantic content from audio signals.

The article is organized into two main parts. In the first part, we highlight how music processing can serve as a tangible and approachable real-world application of signal processing methods. As such, music-based examples aid students in moving from recalling and reciting signal processing concepts (the lowest level of Bloom's taxonomy) toward comprehension and application (the second and third levels). In particular, we give a gentle introduction to Fourier analysis motivated by music and discuss the basic properties of music signals via Fourier analysis.

In the second part, we discuss the role of software tools in signal processing education and detail how they add varying depths of interaction, supporting students' development toward independent work in signal processing. We first explore constrained interaction through the FMP notebooks [7], which closely follow parts of the textbook *Fundamentals of Music Processing* [8] and provide interactive activities to enhance the teaching and learning of classical signal processing techniques. As such, the FMP notebooks carry students further up Bloom's taxonomy, solidifying their abilities to apply signal processing concepts in musical examples. From the instructor's perspective, the inclusion of the FMP notebooks adds interactive elements to the typical lecture-based signal processing classroom with little effort on the instructor's part.

We then discuss incorporating the Python package librosa [9], which enables broader experimentation with signal processing concepts. With the experience gained using the FMP notebooks, students have the technical skills and conceptual fluency to synthesize their signal processing knowledge. Then, adopting and creating programming scripts from the elements in librosa allows students to delve deeper into the implications—in a music context—of altering and exploring the role of various parameters common in signal processing. From the instructor's perspective, including librosa can add depth to examples in a typical lecture-based course as well as provide the ingredients for an end-of-semester project.

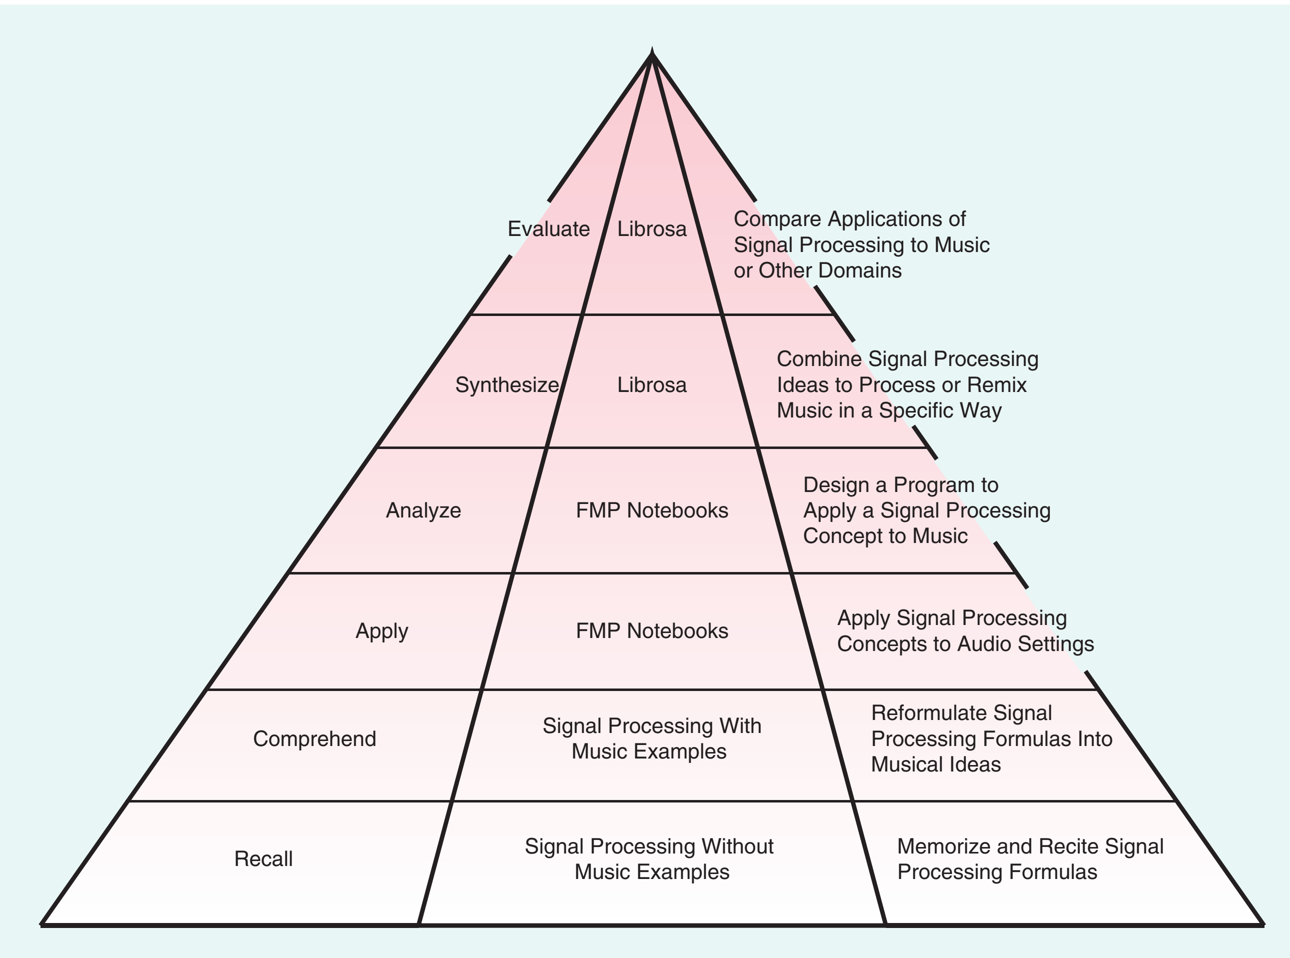

FIGURE 1. An adaptation of Bloom's taxonomy [4] based on levels from Starr et al. [5]. Students begin at the lowest level. Layering music examples, the fundamentals of music processing (FMP) notebooks, and librosa can help them transition to the highest levels of understanding.

The music domain

A typical introductory course on digital signal processing covers a range of topics, including but not limited to: sampling theory, Fourier analysis and the discrete Fourier transform (DFT), convolution and filtering, time–frequency representations, and so on. Although these topics are each fundamental and broadly applicable to a wide array of settings, they can also be conceptually difficult to grasp for newcomers.

When teaching a challenging concept, instructors attempt to find a compelling example that their students can hold onto through the sea of equations and subtleties. For signal processing, music can be that motivating example and help anchor the abstract concepts in a concrete, familiar context. This kind of contextualized pedagogical practice has been shown to improve student retention in the computer science curricula [3]. As such, using music as a context for signal processing gives students an avenue for explaining signal processing in their own words, which moves students from simply recalling formulas and reciting facts (the lowest level of learning on Bloom's taxonomy) to deeper levels of comprehension.

As a multimedia domain, music offers a wide range of data types and formats, including text, symbolic data, audio, image, and video [8], [10]. For example, music can be represented as printed sheet music (often available in the form of digitized images), encoded as Musical Instrument Digital Interface (MIDI) or MusicXML files (structured textual data), and played back as audio recordings. In this article, our primary focus is on music signals or audio representations that encode acoustic waves as generated by an instrument (or voice) and are transmitted through the air as pressure oscillations. As opposed to scores and most symbolic representations, an audio representation encodes the information needed to reproduce a specific acoustic realization of a piece of music. This includes the temporal, dynamic, and tonal microdeviations that make up the particular performance style of a musician. However, in an audio representation, note parameters such as onset times, pitches, or note durations are not given explicitly.

Audio representations of music connect to several standard concepts in signal processing. In Figure 2, we have the musical score, which is performed to create an audio representation. We can then derive several signal processing concepts, including the waveform, spectrogram, and fundamental frequencies. Constructing these different representations and concepts can be a challenging signal processing problem, which, in turn, can provide a motivating example to understand and apply signal processing concepts appropriately.

As a small side note, we want to mention in this context that the design of computational approaches for converting an acoustic music signal into some form of music notation—a task commonly referred to as automatic music transcription—is one of the most challenging and fascinating research problems in signal processing and artificial intelligence. As detailed by Benetos et al. [11], it comprises several subtasks, including multipitch estimation, onset and offset detection, instrument recognition, beat and rhythm tracking, interpretation of expressive timing and dynamics, and score typesetting, to name a few. In the remainder of this section, we explore the connection between music and signal processing and how incorporating music examples into a standard signal processing course provides students with a method for a more nuanced understanding of the topic.

Understanding Fourier analysis through music

Music signals present numerous opportunities to pose natural questions that may be understood by a novice. At very short time scales, one might want to know the fundamental frequency of a played note; at longer time scales, one may ask about the timbre of different instruments; moving to even longer time scales allows us to inquire about higher-level concepts such as melody and rhythm. Each of these questions can be addressed by different

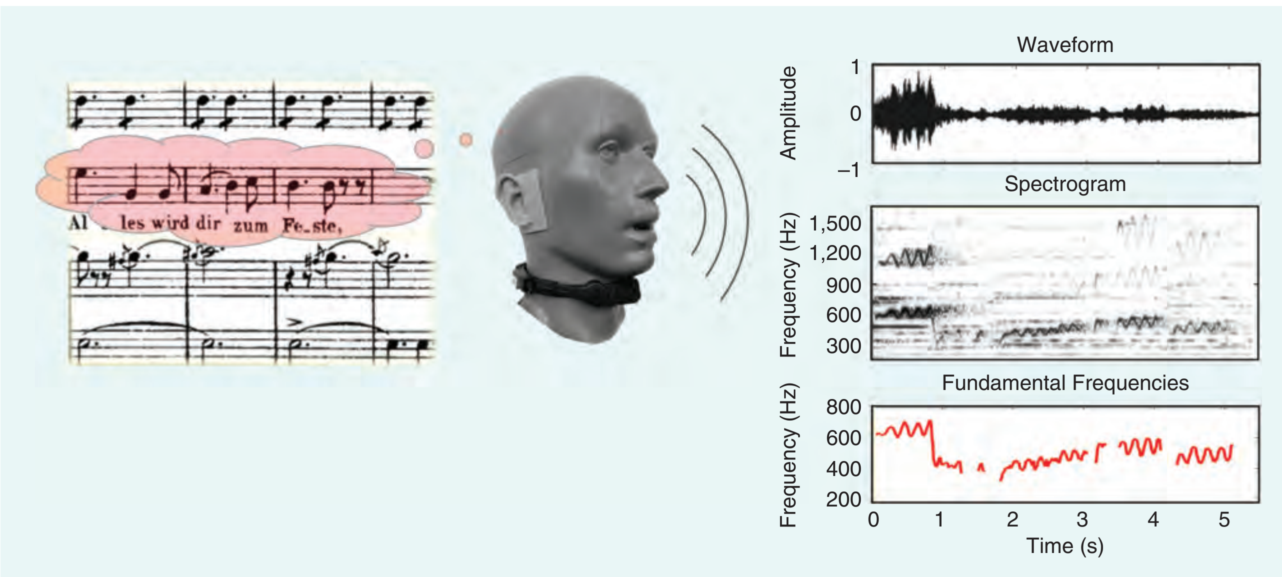

FIGURE 2. A singing scenario, illustrated by a short excerpt of the romantic opera "Der Freischütz" by Carl Maria von Weber. A singer performs a melody accompanied by some instruments. Analyzing the resulting music recording (waveform) leads to challenging signal processing problems (e.g., extracting the singer's fundamental frequency trajectory).

applications of Fourier analysis, and exposure to the same basic principle in multiple related contexts can help solidify a student's understanding. Both the familiarity and complexity of music makes it an ideal motivation and a vehicle for learning signal processing.

The key idea underlying Fourier analysis is to represent arbitrary signals as combinations of sinusoids. When introducing Fourier analysis, it is natural to start with simple combinations consisting of a single sinusoid only, i.e., a pure tone. A sinusoid is completely specified by three parameters: its frequency (the number of oscillations/s, measured in hertz), its amplitude (the peak deviation of the sinusoid from its mean), and its phase (determining where in its cycle the sinusoid is at time zero). Thinking of frequency as the rate of vibration, it is easy to understand that the higher the frequency of a sinusoidal wave, the higher it sounds. This physical analogy yields an intuition for signal processing. For example, a sinusoid having a frequency of 440 Hz (physical attribute) corresponds to the pitch A4 (musical and perceptual attribute). Similarly, the amplitude of a sinusoidal wave relates to the intensity of the sound from a musical instrument.

With this intuition in mind, the aim of Fourier analysis can be interpreted as a kind of reverse engineering problem. Given a music signal, the aim is to measure the intensity with which a sinusoidal wave of a given frequency occurs in (or, more precisely, correlates with) the signal. The collection of intensity values for all frequencies (concealing the role of the phase for the moment) is commonly referred to as Fourier transform. Plotting the intensity values over a frequency axis yields a visualization that reveals the signal's frequency spectrum.

As an example, let us consider the note C4, having a fundamental frequency of 261.6 Hz. When playing this note on different instruments, we hear different sounds. Figure 3 illustrates the waveforms for C4 when played on a piano, trumpet, violin, and flute. Looking at the respective Fourier transform (frequency–intensity plot), one can observe peaks at the fundamental frequency $f = 261$ Hz and its harmonics $2f$, $3f$, $4f$, and so on. However, these plots are not identical. While the peak values drop with increasing frequency for the piano and the violin after $f = 261$ Hz, the highest peak occurs at the fourth harmonic ($4f = 1{,}046$ Hz) for the trumpet and at the second harmonic ($2f = 523$ Hz) for the flute. The distribution of the signal's energy across the harmonics is one important characteristic for the timbre or tone color of an instrument.

The Fourier transform yields frequency information that is integrated over the entire time domain, but most signals are not stationary, and their frequency contents change over time. This observation leads to another central technique in signal processing: short-time Fourier transform (STFT). Instead of considering the entire signal at once, the main idea of the STFT is to partition the signal into small sections in time and consider each smaller section individually. To this end, one fixes a window function, which is a function that is nonzero for only a short period of time (defining the considered section). The original signal is then multiplied with the window function to yield a windowed signal. To obtain frequency information at different time instances, one shifts the window function across the time axis and computes a Fourier transform for each of the resulting windowed signals.

The STFT yields a 2D representation of the original signal that can be visualized by means of a 2D image known as a spectrogram. In this image, the horizontal axis represents time, and the vertical axis represents frequency. Musical signals provide a concrete demonstration of the utility of the STFT: changes in pitch, loudness, or other musically salient characteristics are directly observable in the spectrogram image.

The distinction between the STFT and a standard Fourier transform can often be confusing to students. Additionally, the STFT introduces new parameters not present in the standard Fourier transform, notably the window length (also called the frame length) and the frame rate (as dictated by the hop length or number of samples between successive frames). Without a grounding context, these parameters have no obvious default setting, and students may not immediately grasp the effect of these parameters on the resulting analysis. However, musical signals provide a means for demonstrating the effects of these parameters and, by appealing to basic psychoacoustics, provide a way to connect the values of these parameters to real phenomena. For example, the frame length can be connected to a minimal (perceptible) nontrivial analysis frequency. In this setting, music provides a distinct advantage over other example stimuli (e.g., speech or environmental sound) as it is relatively easy to find (or construct) musical examples that probe the extremal cases of STFT parameters to demonstrate the behavior.

Returning to Figure 3, which shows the spectrograms for note C4, one can observe horizontal lines that are stacked on top of each other for all four instruments. These equally spaced lines correspond to partials, which are sinusoidal sound components that are not necessarily—but are often close to—harmonics. In the case of the piano, the higher partials contain less and less of the signal's energy. Furthermore, the decay of a piano sound over time is reflected by the fading out of the horizontal lines. For the trumpet sound, the spectrogram indicates that the signal's energy is concentrated more in the higher partials. Also, opposed to the piano sound, there does not seem to be any intensity decay over time, indicating that the trumpet player keeps the volume of the sound constant.

While this is also the case for the violin and flute sounds, one can observe other phenomena that typically go along with certain playing styles, such as vibrato. For example, when looking at the waveform, one can observe periodic variations in amplitude, also referred to as amplitude modulation. In the spectrogram (particularly visible in the flute example), these variations appear as the regular pulsation of intensity values along the time dimension. Amplitude modulations often go along with frequency modulations, which are regular, pulsating changes of frequency over time. In the spectrogram (particularly visible in the violin example), these modulations appear as wave-like oscillations along the time dimension.

Both amplitude and frequency depend on two parameters: the extent of the variation and the rate at which the amplitude or frequency is varied. Even though they are simply local changes in intensity and frequency, the modulations do not necessarily evoke a perceived change in loudness or pitch of the overall musical tone. Rather, they are features that are used by musicians to influence the timbre of a musical tone.

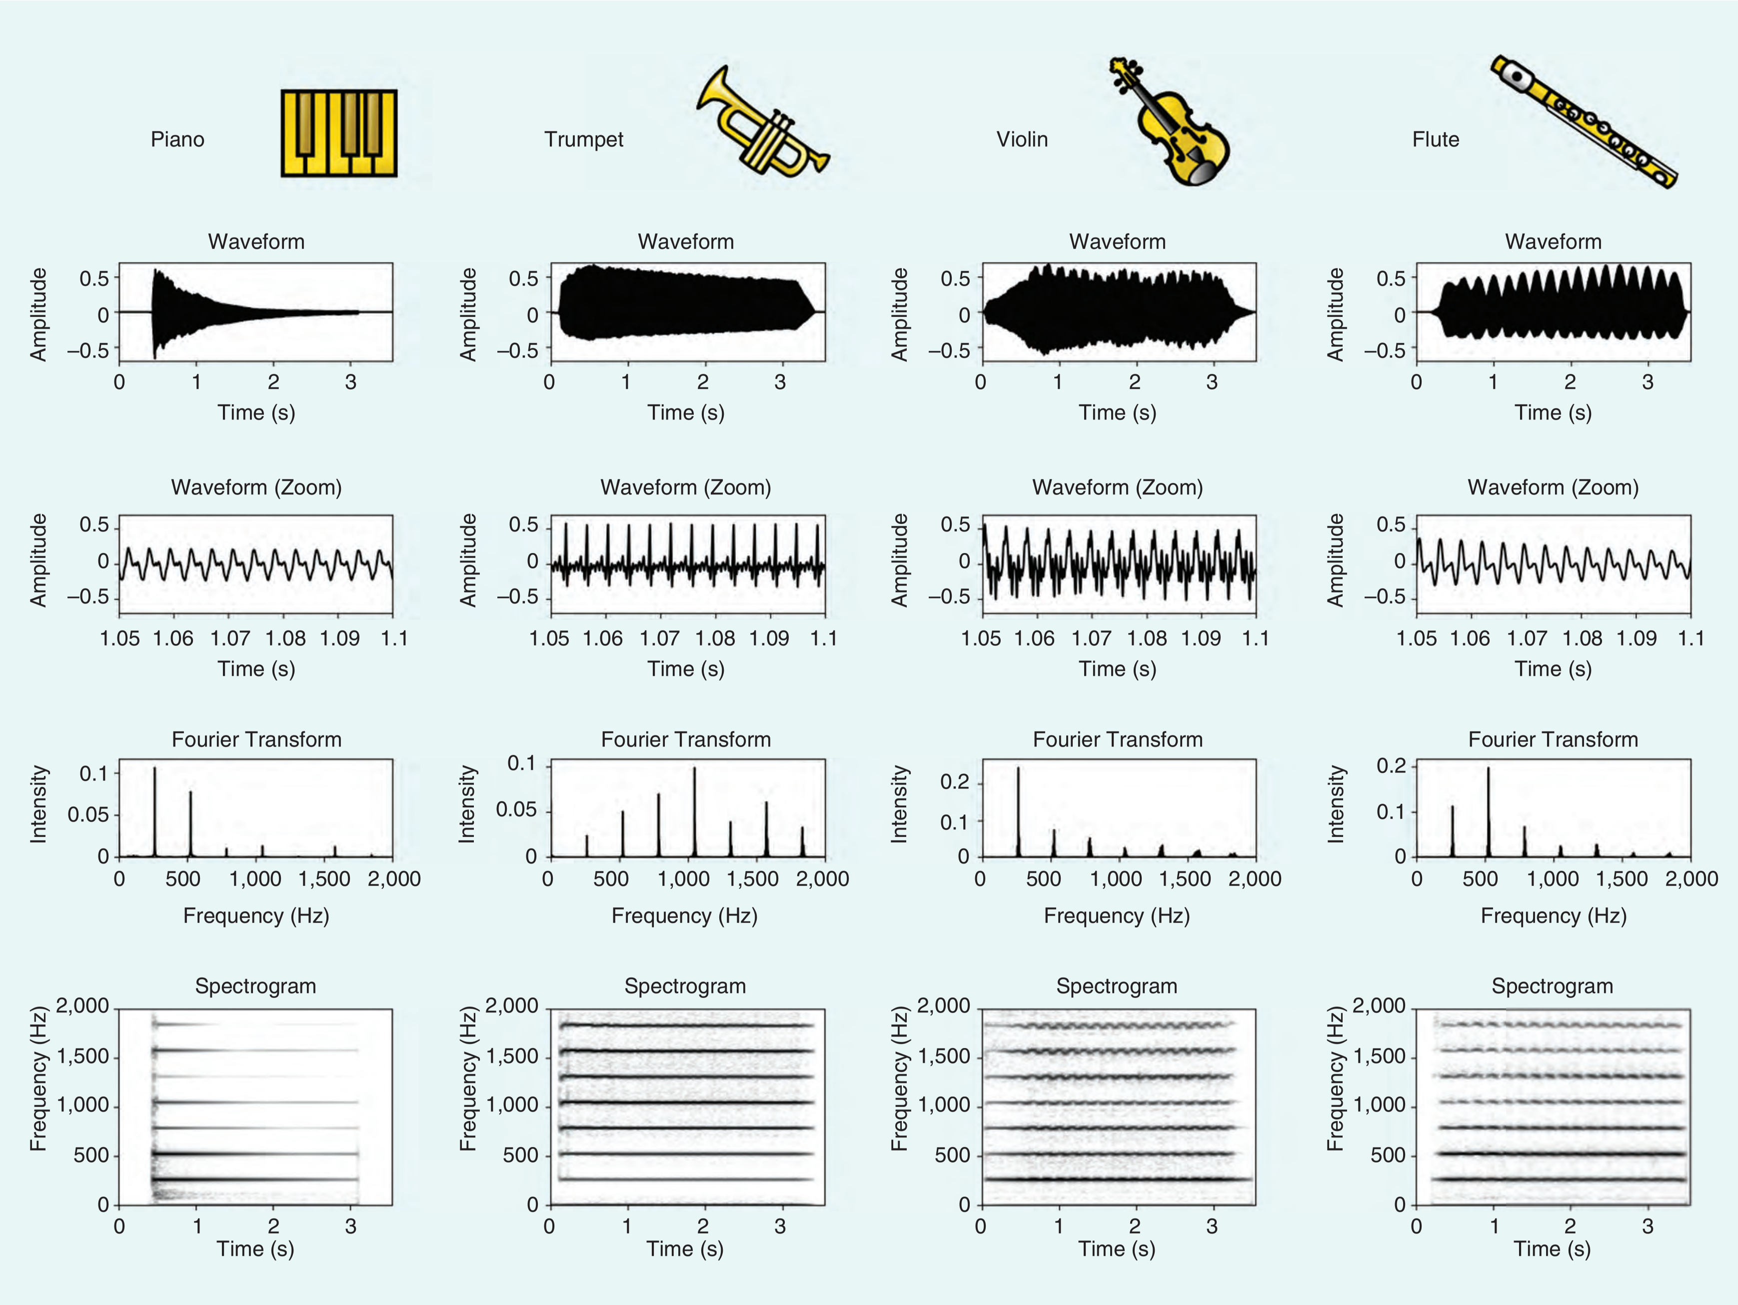

FIGURE 3. The waveform, Fourier transform, and spectrogram for different instruments playing the same note C4 (fundamental frequency 261.6 Hz).

The practical application of Fourier analysis to music

So far, we have established that music can provide intuition for the Fourier transform and its short-time version, and we have demonstrated how the Fourier transform and STFT can provide an intuitive understanding of a music signal's properties. With this connection between music and signal processing established, music scenarios such as singing (see Figure 2) can motivate students to explore the potential of Fourier analysis in an interactive and playful fashion. Most students and schoolchildren are familiar with music video games (e.g., "SingStar" or "Rock Band"), where the task is to sing along with the music to score points. To compare the singer's input waveform with the game's reference melody, one could employ Fourier analysis. This makes it possible to convert the waveform into a sequence (or trajectory) of fundamental frequency values.

Such trajectories are often made visible in the video games, superimposed with piano roll-like visualizations of reference pitches. One can mimic the basic idea of such games by employing real-time-capable software for visualizing the frequency content of sounds while singing. Deepening the understanding of signal properties, such software also allows students to experiment with algorithmic parameters that control the STFT's time and frequency resolution as well as the intensity visualization (e.g., switching from a linear to a decibel scale).

Besides analyzing the melodic properties of music signals, the Fourier transform can be applied for many more music processing tasks, including harmony analysis, instrument recognition, rhythmic analysis, and source separation. In the following, we examine beat tracking. Temporal and structural regularities are perhaps the most important incentives for people to get involved and interact with music [8], [12]. It is the beat that drives music forward and provides the temporal framework for a piece of music. Intuitively, the beat corresponds to the pulse that a human taps along with when listening to music [13]. The term tempo [often specified in beats per minute (BPM)] refers to the rate of the pulse and is given by the reciprocal of the beat period.

The beat tracking task seeks to extract the beat and tempo information from audio recordings, and it is one of the central and most well-studied research problems in MIR. Beat tracking is an instructive, challenging, and multifaceted application for teaching and learning signal processing. Most approaches to beat tracking are based on two assumptions: first, the beat positions correspond to note onsets (often percussive in nature), and second, beats are periodically spaced in time. Note that for certain types of music these two assumptions may be questionable. For example, in passages with syncopation, beat positions may not go along with any onsets, or the periodicity assumption may be violated for romantic piano music with strong tempo fluctuations (played rubato). The explicit modeling of such simplifying assumptions is at the core of researching and teaching music processing.

The two assumptions to beat tracking motivate approaching it in two steps. Consider the short excerpt of "Another One Bites the Dust" by Queen, depicted in Figure 4, for a concrete visualization as we examine each of these steps. In the first step, one often estimates the positions of starting times of notes or other musical events as they occur in a music signal—a task commonly referred to as onset detection. As illustrated in Figure 4, onsets often go along with a sudden change in a signal's properties. Such changes may be seen as sharp amplitude increases in the waveform. For notes with soft onsets or complex music with several instruments playing at the same time, the detection of individual note onsets becomes much harder. In these cases, converting the signal into a spectrogram turns out beneficial. In particular, percussive onsets, as produced by a drum or hi-hat, result in vertical lines in the spectrogram.

This phenomenon comes from Fourier analysis: the energy of transient events is spread across the entire spectrum of frequencies, thus yielding broadband spectral structures. To detect these structures, one basic idea is to compute a kind of distance between subsequent column vectors of the spectrogram. This results in a novelty function (also known as the spectral flux) that captures the sudden changes in the signal's frequency distribution. The peaks of such a novelty function are good indicators for note onset candidates.

In the second step, based on the assumption that beats are periodically spaced in time, the novelty function is analyzed with regard to reoccurring patterns. This is a particularly suitable, intuitive, and concrete setting for studying different techniques of periodicity analysis—a central concern of signal processing and time series analysis. One approach based on autocorrelation analysis aims to detect periodic self-similarities by comparing a novelty function with time-shifted copies [14]. An alternative approach uses a bank of comb filter resonators, where a novelty function is compared with templates consisting of equally spaced spikes, with each template representing a specific tempo [15]. A third approach compares the novelty function with sinusoidal templates, each corresponding to a specific frequency. This is exactly the idea of Fourier analysis, yielding a frequency representation of the novelty function.

Starting with playing and listening to music, a teacher can smoothly transition to introducing the basic concepts of signal processing. Singing analysis and beat tracking are two tangible example tasks that help transition learning signal processing from one of rote memorization to a contextually meaningful pursuit. Additional examples can be found in the *Fundamentals of Music Processing* textbook [8]. Although our focus in this section has been on the applications of basic Fourier analysis, many music analysis tasks can be easily adapted to more advanced topics, such as wavelet theory [16]. In the next section, we discuss educational software tools for teaching and learning such concepts.

Educational software tools

In addition to motivating and tangible music-based scenarios, the availability of suitably designed software packages that make signal processing more accessible is crucial in view of interactive learning [17]. Over the last 20 years, as MIR developed as a research field, so did computational accessibility, and the MIR community has contributed with several excellent toolboxes that provide modular source code for processing and analyzing music signals. Prominent examples are essentia [18], madmom [19], Marsyas [20], and the MIRtoolbox [21]. These toolboxes are mainly designed for research-oriented access to audio processing, yielding code for audio feature extraction as well as for various MIR applications. Here, we focus on two concrete software examples: the FMP notebooks [7] that have an explicit educational lens and the Python package librosa [9] that has become a standard in MIR research—and recently has also been incorporated into introductory MIR courses. We describe how these tools facilitate multiple entry points to delve deeper into classical signal processing techniques while bridging the gap between education and cutting-edge research.

FMP notebooks

FMP notebooks offer an interactive foundation for MIR and for teaching and learning FMP [7], which, when used in a traditional signal processing course, can enhance students' understanding of the topic. By closely following the eight chapters of the textbook [8], the FMP notebooks provide an explicit link between structured educational environments and current professional practices in line with current curricular recommendations for computer science [22]. Furthermore, these notebooks provide a vehicle for students to transition between comprehending signal processing ideas in their own words (the second level of Bloom's taxonomy) toward applying these ideas interactively to music examples (the third and fourth levels of Bloom's taxonomy). For many MIR tasks, fundamental algorithms and signal processing techniques are discussed in detail. An overview of the main topics covered by the FMP notebooks is presented in Figure 5. Besides the treatment of the theory, the notebooks demonstrate how these techniques can be implemented by providing specific Python code examples.

The FMP notebooks leverage the Jupyter notebook framework [23], which has become a standard in industry as well as in educational settings. This open source web application allows users to create documents that contain live code, text-based information, mathematical formulas, plots, images, sound examples, and videos. Jupyter notebooks are often used as a publishing format for reproducible computational workflows [23]. They can be exported to a static HTML format, which makes it possible to generate web applications that can be accessed through standard web browsers with no specific technical requirements. By leveraging the Jupyter framework, the FMP notebooks bridge the gap between theory and practice by interleaving technical concepts, mathematical details, code examples, illustrations, and sound examples within the unifying Jupyter framework (see Figure 6). Additionally, the notebooks are essentially self-contained in terms of content by including introductions for each MIR task, providing important mathematical definitions, and describing the computational approaches in detail.

One primary purpose of the FMP notebooks is to provide audiovisual material as well as Python code examples that implement the computational approaches step by step. Additionally, the FMP notebooks provide an interactive framework that allows students to experiment with their own music examples, explore the effect of parameter settings, and gain an understanding of the computed results by suitable visualizations and sonifications. These functionalities are examples of "procedural literacy" [24], centering theoretical discussions around computational procedures.

The FMP notebooks, even though they contain a library of MIR functions (called libfmp), are not designed to be a toolbox per se. Instead, for a given music processing pipeline, the FMP notebooks introduce the code in a step-by-step fashion interleaved with explanations. This allows a student to access, visualize, and understand the intermediate steps. We illustrate this principle by coming back to our beat tracking scenario. In Figure 4, we introduced a spectrum-based novelty function, the peaks of which indicate note onset candidates.

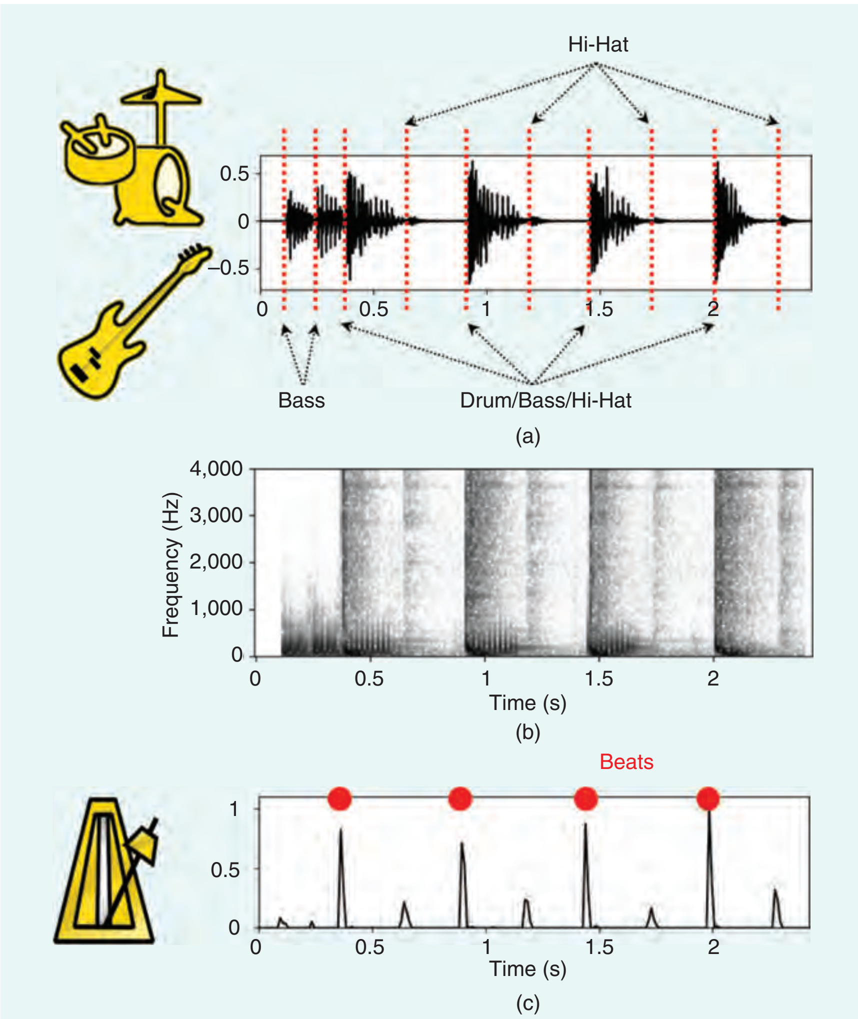

FIGURE 4. A short excerpt of "Another One Bites the Dust" by Queen. The figure shows (a) the waveform with annotated onset positions, (b) a spectrogram, and (c) a novelty function with annotated beat positions.

We now apply the same concept to an orchestral recording of Waltz No. 2 by Dimitri Shostakovich's Suite for Variety Orchestra No. 1. Figure 7 illustrates the score (in a piano-reduced version) as well as the novelty function of a short excerpt of this piece. Note that the first beats (downbeats) of the 3/4 meter are played softly by nonpercussive instruments, leading to relatively weak and blurred onsets. In contrast, the second and third beats are played sharply ("staccato"), supported by percussive instruments. These properties are also reflected by the spectral-based novelty function: the peaks that correspond to downbeats are hardly visible or even missing, whereas the peaks that correspond to the percussive beats are much more pronounced.

As for beat tracking, we again use Fourier analysis—this time applied to the novelty function rather than to the audio signal. As for the STFT, the idea is to locally compare a given novelty function with windowed sinusoids. This time, the frequency of the sinusoid is interpreted in terms of BPM (e.g., an oscillation rate of 1 Hz corresponds to 60 BPM). The resulting spectrogram is then called a tempogram, where the frequency axis is interpreted as a tempo axis. Besides the frequency, we also use the phase information of the complex STFT coefficients to determine for each time position a windowed sinusoid that best captures the local peak structure of the novelty function. This is illustrated in Figure 7.

Instead of looking at the windowed sinusoids individually, the idea is to employ an overlap-add technique by accumulating all of the locally optimal sinusoids over time. As a result, one obtains a single function that can be regarded as a local periodicity enhancement of the original novelty function. Revealing PLP information, this representation is referred to as a PLP function [25]. Having a pronounced peak structure, the beat positions can now be obtained from the PLP function using a simple peak picking strategy.

By looking at this concrete example, we illustrated how the FMP notebooks yield explicit access to all of the intermediate steps, starting with a musical score and ending with a sonification of the detected beat positions superimposed with the original audio recording. Furthermore, this example showed how Fourier analysis could be applied for periodicity enhancement while highlighting the role of the phase. When teaching and learning signal processing, we advocate that it is essential to have a holistic view of the MIR task at hand, the algorithmic approach, and its practical implementation. Looking at all of the steps of the processing pipeline sheds light on the input data and its biases, possible violations of model assumptions, and the shortcomings of quantitative evaluation measures. Only by an interactive examination of all of these aspects will students acquire a deeper understanding of the concepts, transitioning from merely explaining concepts (the lowest level of Bloom's taxonomy) to applying their signal processing approaches both conceptually and in code.

Furthermore, the imprecise definitions of the MIR tasks allow for a richer discussion and cognitive interaction with signal processing and music. For example, in beat tracking, there are a number of questions that naturally arise: What is an onset? Can it be described by a single

Part	Title	Notions, Techniques, and Algorithms
B 	Basics	Basic Information on Python, Jupyter Notebooks, Anaconda Package Management System, Python Environments, Visualizations, and Other Topics
0 	Overview	Overview of the Notebooks (https://www.audiolabs-erlangen.de/FMP)
1 	Music Representations	Music Notation, MIDI, Audio Signal, Waveform, Pitch, Loudness, Timbre
2 	Fourier Analysis of Signals	Discrete/Analog Signal, Sinusoid, Exponential, Fourier Transform, Fourier Representation, DFT, FFT, STFT
3 	Music Synchronization	Chroma Feature, Dynamic Programming, Dynamic Time Warping, Alignment, User Interface
4 	Music Structure Analysis	Similarity Matrix, Repetition, Thumbnail, Homogeneity, Novelty, Evaluation, Precision, Recall, F-Measure, Visualization, Scape Plot
5 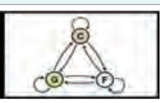	Chord Recognition	Harmony, Music Theory, Chords, Scales, Templates, Hidden Markov Model, Evaluation
6 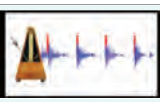	Tempo and Beat Tracking	Onset, Novelty, Tempo, Tempogram, Beat, Periodicity, Fourier Analysis, Autocorrelation
7 	Content-Based Audio Retrieval	Identification, Fingerprint, Indexing, Inverted List, Matching, Version, Cover Song
8 	Musically Informed Audio Decomposition	Harmonic/Percussive Separation, Signal Reconstruction, Instantaneous Frequency, Fundamental Frequency, Trajectory, Nonnegative Matrix Factorization

FIGURE 5. An overview of the main topics covered by the FMP notebooks (adapted from [7] and [8]). FFT: fast Fourier transform.

time instant? Does beat tracking make sense for certain musical passages (e.g., music with rubato)? Would humans agree when asked to specify a single tempo value? Is the evaluation metric relevant for a given application? In a curriculum on signal processing, wrestling with such questions illuminates to students the challenges of computational approaches in the applied sciences.

The librosa python package

The FMP notebooks provide an interactive framework in which students can learn about and experiment with signal processing and MIR algorithms. However, when students transition from learning to professional practice and research, we expect them to outgrow the FMP notebooks and begin developing their own digital signal processing methods and programs, corresponding to students' arrival at the top of Bloom's taxonomy. However, this transition can be difficult without proper infrastructural (i.e., software) support. The librosa package was designed to fill this role, providing standardized and flexible reference implementations of many common methods in MIR [9].

Whereas the FMP notebooks are designed to introduce the fundamental concepts in signal processing, librosa is intended to facilitate the high-level composition of basic methods into complex pipelines. As its original intended audience was the MIR research community, it was designed to facilitate the development of experimental research code. In the seven years since its first release, numerous scholarly publications have used librosa to provide the underlying signal processing framework, and many of these publications include open source software that students and researchers can download, use, and extend the work. The availability of open research software provides an avenue to train new researchers as they can directly see how prior work was done and have significantly less work to do if modifying or extending it to achieve a new goal.

Conversely, librosa itself provides reference implementations of many previously published methods (various feature extractors, phase retrieval methods, spectrogram decompositions, beat tracking algorithms, and so on) that can be used to independently replicate a method with relatively little effort. To demonstrate this, a collection of advanced examples are provided in the documentation, several of which demonstrate how to fully reproduce a published method by combining the building blocks provided by librosa. These examples can also be exported as Jupyter notebooks that a user can download and run on their own machine. This feature of the documentation serves a similar, interaction-oriented goal as the FMP notebooks, except that it is aimed principally at researchers and software developers already familiar with the fundamentals.

The design of librosa

The application programming interface (API) of librosa was intentionally designed to present a low barrier to entry for new users, eschewing complex class hierarchies and object-oriented interfaces in favor of simple data structures (numerical arrays) and functions. This kind of function-oriented design can be easier for new users to learn as functions (in contrast to objects) have well-defined entry and exit points and no internal state for the programmer to understand (i.e., all parameters are explicitly visible in the call signature). Similarly, variable names are consistently defined across the package and are human readable (e.g., n_fft

FIGURE 6. An overview of didactical aspects of the FMP notebooks and their implementation using the interactive Jupyter notebook framework.

instead of N for the number of analysis frequencies in a Fourier transform). (These design principles, among many others found in the scientific Python community, were clearly articulated by Gaël Varoquaux's 2017 keynote address at the annual SciPy conference [26].)

Figure 8 provides a brief example code listing that follows the rhythm analysis example given in Figure 7. While the library allows a user to directly construct the PLP function from an audio signal (line 7), a user can also explicitly construct intermediate representations, such as the novelty function, directly (lines 10 and 11). The resulting code is still compact and high level, but it also facilitates exploration and experimentation. For example, a user can easily change the calculation of the novelty function and leave the remainder of the PLP analysis fixed, allowing them to carefully measure the result of their interventions. This design philosophy is not limited to this one example but rather is seen throughout the library as a whole.

Beyond its API design, the library developers strive for complete and thorough documentation, fully worked code examples for each function, and well-documented and readable source code. The latter point is enforced by stringent code review and ensures a high standard of quality for all source code contributions. Well-written source code can be instructional, both in demonstrating clearly how any particular method works and in providing more general examples of how to structure complex programs and libraries. The intent behind the emphasis on code quality is to allow any knowledgeable user to look into the code and, with minimal effort, quickly be able to understand how it works—and potentially extend it with new contributions. More generally, the library itself provides an example of many

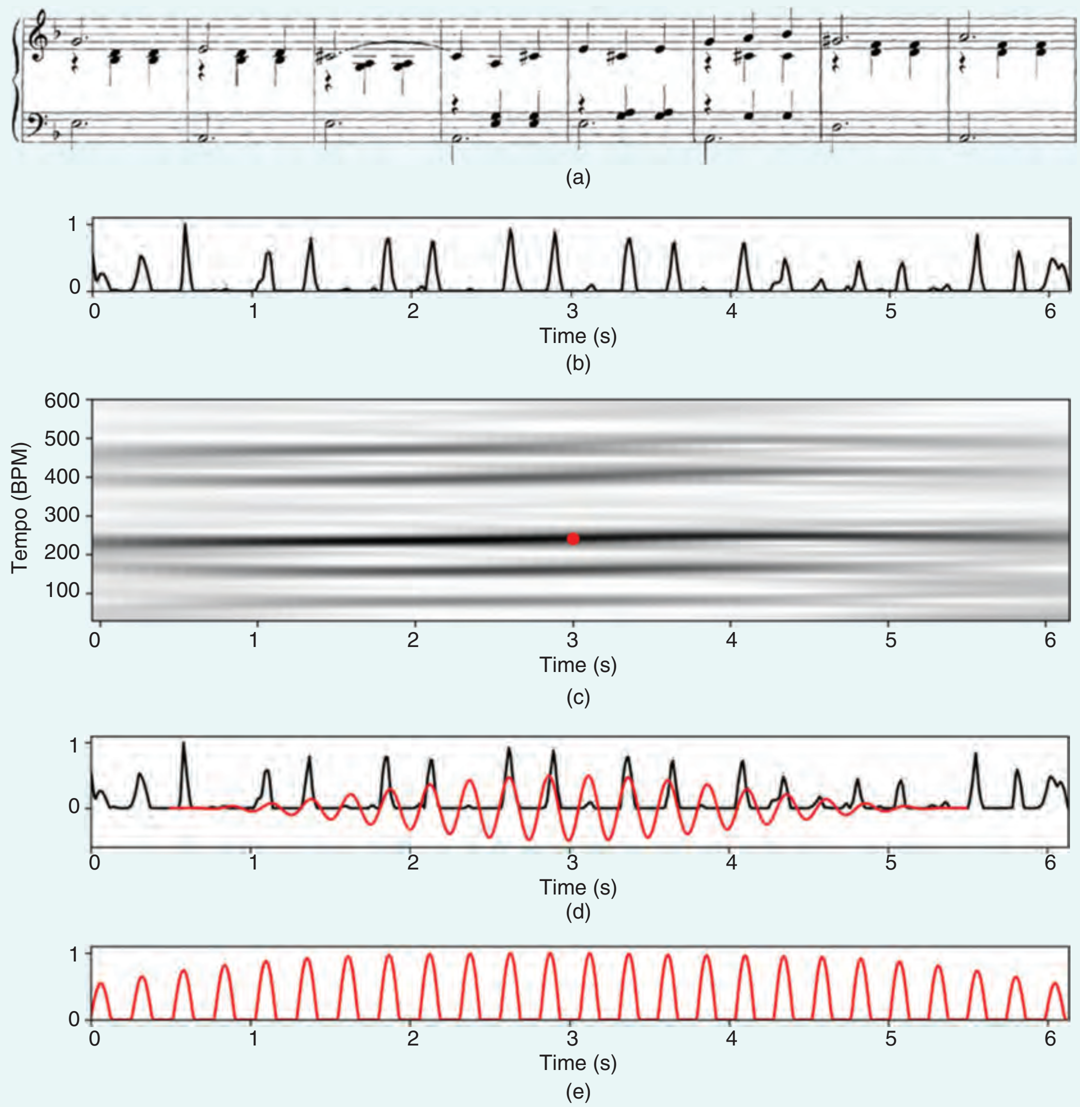

FIGURE 7. An illustration of the music processing pipeline for computing the predominant local pulse (PLP) function given a novelty function (see also Figure 4). (a) A small excerpt of Waltz No. 2 by Dimitri Shostakovich's Suite for Variety Orchestra No. 1. (b) The novelty function, (c) the Fourier tempogram, (d) the novelty function with a windowed sinusoid, and (e) the PLP function.

software engineering best practices that can be readily adopted in research settings, such as version control, code review processes, and continuous integration testing [27].

Both the FMP notebooks and librosa create opportunities for interaction in a traditional signal processing course. Adding just the FMP notebooks adds interactions similar to a "coding worksheet" that allows students to dynamically play with examples. Additionally, students become familiar with the Jupyter framework and Python syntax. Including librosa allows for broader experimentation with signal processing and provides a common language for cumulative course projects, where students can demonstrate their familiarity and creativity with a number of signal processing topics. Furthermore, leveraging tools like the FMP notebooks and librosa provides opportunities for learning beyond signal processing. By employing these Python- and Jupyter-based tools, students are passively learning about open access and reproducibility, two topics that cut across disciplines.

```
import librosa

# Load an audio file, returns the signal and sampling rate
y, sr = librosa.load('Shostakovich_Waltz.wv')

# Compute the PLP pulse directly from the signal
pulse = librosa.beat.plp(y=y, sr=sr)

# Or we can use a pre-computed novelty function
nov = librosa.onset.onset_strength (y=y, sr=sr)
pulse = librosa.beat.plp(onset_envelope=nov, sr=sr)

# Compute the Fourier tempogram
tgram = librosa.featue.fourier_tempogram(y=y, sr=sr)
```

FIGURE 8. An example usage of librosa to reproduce the analysis illustrated in Figure 7. All of the results are stored as numerical arrays and can be directly manipulated by the user. The functionality is grouped into different submodules (onset, feature, beat), and the interfaces for high-level analyses (e.g., PLP) support exploration by allowing users to precompute intermediate representations.

Interactivity and learning

In this article, we have presented a scaffold to naturally transition students from starting to learn about signal processing to beginning independent research. Using music examples and technologies provide what Guzdial [3] calls a contextualized educational experience for signal processing. Concretely, we have demonstrated how music-based examples can make Fourier analysis more tangible to beginners. We then demonstrated how the FMP notebooks provide a constrained scaffold for interacting with music examples for signal processing. Building on the skills and conceptual understanding gained through concrete and interactive music-based examples, students can explore more advanced applications of Fourier analysis through broader experimentation via the Python library librosa.

The incorporation of music into a signal processing course provides students an avenue for "inauthentic legitimate peripheral participation" [28] in the field of signal processing. Extending the work of Lave and Wenger [6] by defining "legitimate peripheral participation" as building a scaffold for introducing students to concepts similar to the apprentice structure, Guzdial and Tew [28] connect this model to computing courses that have media-based examples (such as photographs) as the basis for the coursework. In this article, we have modeled a similar extension, applying their framework to the specifics of music as a vehicle for learning signal processing by first engaging the highly structured FMP notebooks and then leveraging the range of tools in the librosa package. In effect, this structure brings students through Bloom's taxonomy, helping them to reach the deepest understanding of signal processing concepts.

Signal processing is about finding structure in signals. Audio is a familiar signal modality, and music is explicitly and intentionally structured audio. Leveraging the familiarity of music, Fourier analysis becomes more concrete, and examples from textbooks such as [1] and [8] can be easily incorporated into a traditional signal processing course. For instructors seeking to transition to an interactive alternative to the lecture-based classroom model with a "sage on a stage" simply depositing knowledge into students' brains, the FMP notebooks provide one such alternative. Created in the Jupyter framework with an explicit connection to the *Fundamentals of Music Processing* textbook [8], the FMP notebooks provide an interactive environment where students are invited to grapple with concepts through small structured coding examples. The Jupyter framework underlying the FMP notebooks provides an experimental playground for students to test signal processing concepts on music and manipulate music using signal processing ideas. Once students are familiar with introductory signal processing and MIR concepts, they can continue experimenting via the examples in librosa presented as coding notebooks.

In this article, we leveraged the inherent familiarity of music to motivate theoretical signal processing and extend these examples to an interactive learning experience through the FMP notebooks and librosa. We have discussed how the interplay between music and signal processing leads to a variety of interactions: interaction through applications, hands-on interaction with the material through experimentation, and interaction between the structure of a classroom and the experimentation of research. We have provided resources that instructors can use in their classrooms, and we have been diligent in describing how these resources can be implemented in terms of course activities, from enhancing lectures to incorporating cumulative class projects.

Acknowledgments

The International Audio Laboratories Erlangen is a joint institution of the Friedrich-Alexander-Universität Erlangen-Nürnberg and the Fraunhofer-Institut für Integrierte Schaltungen IIS. Meinard Müller thanks the German Research Foundation (DFG) for various research grants that allowed him to conduct fundamental research in music processing. Katherine M. Kinnaird is the Clare Boothe Luce Assistant Professor of Computer Science and Statistical and Data Sciences at Smith College and, as such, is supported by the Henry Luce Foundation's Clare Boothe Luce Program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Luce Foundation. Finally, we thank Roger Dannenberg for his helpful comments.

Authors

Meinard Müller (meinard.mueller@audiolabs-erlangen.de) received his Ph.D. degree in computer science from the University of Bonn, Germany in 2001. Since 2012, he has held a professorship for Semantic Audio Signal Processing at the International Audio Laboratories Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, 91058, Germany. His research interests include music processing, music information retrieval, audio signal processing, and motion processing. He wrote a monograph *Information Retrieval for Music and Motion* (Springer-Verlag, 2007) and a textbook *Fundamentals of Music Processing* (Springer-Verlag, 2015, www.music-processing.de). In 2020, he was elevated to IEEE Fellow for his contributions to music signal processing.

Brian McFee (brian.mcfee@nyu.edu) received his B.S. degree in 2003 in computer science from the University of California, Santa Cruz, and his M.S. and Ph.D. degrees in 2008 and 2012 in computer science and engineering from the University of California, San Diego. Currently, he is an assistant professor of music technology and data science at New York University, New York, 10003, USA. His work lies at the intersection of machine learning and audio analysis. He is an active open source software developer and the principal maintainer of the librosa package for audio analysis.

Katherine M. Kinnaird (kkinnaird@smith.edu) received her A.B. degree in 2008 in mathematics from Wellesley College and her A.M. and her Ph.D. degrees in 2010 and 2014 in mathematics from Dartmouth College. Currently, she is the Clare Boothe Luce Assistant Professor of Computer Science and Statistical and Data Sciences at Smith College, Northampton, Massachusetts, 01063, USA. Her research takes a mathematical lens to machine learning approaches on cultural artifacts including music and TED talk transcripts. Throughout her career, she has participated in professional development opportunities focused on excellence in teaching, including Dartmouth College's National Science Foundation Graduate Teaching Fellows in K-12 Education program and the Mathematical Association of America's Project NExT (New Experiences in Teaching).

References

[1] A. Lerch, *An Introduction to Audio Content Analysis: Applications in Signal Processing and Music Informatics*. Hoboken, NJ: Wiley, 2012.

[2] S. V. Vaseghi, *Multimedia Signal Processing: Theory and Applications in Speech, Music and Communications*. Hoboken, NJ: Wiley, 2007.

[3] M. Guzdial, "Does contextualized computing education help?" *ACM Inroads*, vol. 1, no. 4, pp. 4–6, 2010. doi: 10.1145/1869746.1869747.

[4] B. S. Bloom and M. D. Engelhart, *Taxonomy of Educational Objectives: The Classification of Educational Goals. Handbook I: Cognitive Domain*. New York: David McKay Co., 1956.

[5] C. Starr, B. Manaris, and R. H. Stalvey, "Bloom's taxonomy revisited: Specifying assessable learning objectives in computer science," *SIGCSE Bull.*, vol. 40, no. 1, pp. 261–265, 2008. doi: 10.1145/1352322.1352227.

[6] J. Lave and E. Wenger, *Situated Learning: Legitimate Peripheral Participation*. Cambridge, U.K.: Cambridge Univ. Press, 1991.

[7] M. Müller and F. Zalkow, "FMP Notebooks: Educational material for teaching and learning fundamentals of music processing," in *Proc. Int. Soc. Music Inf. Retrieval Conf. (ISMIR)*, Delft, The Netherlands, 2019, pp. 573–580.

[8] M. Müller, *Fundamentals of Music Processing*. Berlin: Springer-Verlag, 2015.

[9] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "Librosa: Audio and music signal analysis in python," in *Proc. Python Sci. Conf.*, Austin, TX, 2015, pp. 18–25. doi: 10.25080/Majora-7b98e3ed-003.

[10] G. Wiggins, D. Müllensiefen, and M. Pearce, "On the non-existence of music: Why music theory is a figment of the imagination," *Musicae Scientiae, Discussion Forum*, vol. 14, no. 1, pp. 231–255, 2010. doi: 10.1177/10298649100140S110.

[11] E. Benetos, S. Dixon, Z. Duan, and S. Ewert, "Automatic music transcription: An overview," *IEEE Signal Process. Mag.*, vol. 36, no. 1, pp. 20–30, 2019. doi: 10.1109/MSP.2018.2869928.

[12] R. Parncutt, "A perceptual model of pulse salience and metrical accent in musical rhythms," *Music Perception*, vol. 11, no. 4, pp. 409–464, 1994. doi: 10.2307/40285633.

[13] F. Lerdahl and R. Jackendoff, *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press, 1983.

[14] M. E. P. Davies and M. D. Plumbley, "Context-dependent beat tracking of musical audio," *IEEE Trans. Audio, Speech, Language Process. (2006–2013)*, vol. 15, no. 3, pp. 1009–1020, 2007. doi: 10.1109/TASL.2006.885257.

[15] A. P. Klapuri, A. J. Eronen, and J. Astola, "Analysis of the meter of acoustic musical signals," *IEEE Trans. Audio, Speech, Language Process. (2006–2013)*, vol. 14, no. 1, pp. 342–355, 2006. doi: 10.1109/TSA.2005.854090.

[16] S. Mallat, *A Wavelet Tour of Signal Processing – The Sparse Way*, 3rd ed. New York: Academic, 2009.

[17] M. Guzdial, "Exploring hypotheses about media computation," in *Proc. Int. Comput. Educ. Res. Conf. (ICER)*, La Jolla, CA, 2013, pp. 19–26. doi: 10.1145/2493394.2493397.

[18] D. Bogdanov, N. Wack, E. Gómez, S. Gulati, P. Herrera, O. Mayor, G. Roma, J. Salamon et al., "Essentia: An audio analysis library for music information retrieval," in *Proc. Int. Soc. Music Inf. Retrieval Conf. (ISMIR)*, Curitiba, Brazil, 2013, pp. 493–498.

[19] S. Böck, F. Korzeniowski, J. Schlüter, F. Krebs, and G. Widmer, "madmom: A new Python audio and music signal processing library," in *Proc. ACM Int. Conf. Multimedia (ACM-MM)*, Amsterdam, The Netherlands, 2016, pp. 1174–1178.

[20] G. Tzanetakis, "Music analysis, retrieval and synthesis of audio signals MARSYAS," in *Proc. ACM Int. Conf. Multimedia (ACM-MM)*, Vancouver, British Columbia, 2009, pp. 931–932.

[21] O. Lartillot and P. Toiviainen, "MIR in MATLAB (II): A toolbox for musical feature extraction from audio," in *Proc. Int. Soc. Music Inf. Retrieval Conf. (ISMIR)*, Vienna, Austria, 2007, pp. 127–130.

[22] Joint Task Force on Computing Curricula of the Association for Computing Machinery (ACM) and IEEE Computer Society, *Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science*. New York: Association for Computing Machinery, 2013.

[23] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick et al., "Jupyter notebooks—A publishing format for reproducible computational workflows," in *Proc. Int. Conf. Electron. Publishing*, Göttingen, Germany, 2016, pp. 87–90.

[24] A. Vee, "Understanding computer programming as a literacy," *Literacy Composition Stud.*, vol. 1, no. 2, 2013. doi: 10.21623/1.1.2.4.

[25] P. Grosche and M. Müller, "Extracting predominant local pulse information from music recordings," *IEEE Trans. Audio, Speech, Language Process. (2006–2013)*, vol. 19, no. 6, pp. 1688–1701, 2011. doi: 10.1109/TASL.2010.2096216.

[26] G. Varoquaux, "Coding for science and innovation," Austin, TX, 2017. Accessed: Jan. 29, 2021. [Online]. Available: https://pyvideo.org/scipy-2017/keynote-coding-for-science-and-innovation.html

[27] B. McFee, J. W. Kim, M. Cartwright, J. Salamon, R. M. Bittner, and J. P. Bello, "Open-source practices for music signal processing research: Recommendations for transparent, sustainable, and reproducible audio research," *IEEE Signal Process. Mag.*, vol. 36, no. 1, pp. 128–137, 2019. doi: 10.1109/MSP.2018.2875349.

[28] M. Guzdial and A. E. Tew, "Imagineering inauthentic legitimate peripheral participation: An instructional design approach for motivating computing education," in *Proc. Int. Comput. Educ. Res. Workshop (ICER)*, 2006, pp. 51–58.

SP

Kathleen E. Wage, John R. Buck,
Jill K. Nelson, and Margret A. Hjalmarson

What Were They Thinking?

Refining conceptual assessments using think-aloud problem solving

©SHUTTERSTOCK.COM/LUCKYSTEP

As today's practitioners and educators, we bear responsibility for training our students to achieve the deep conceptual understanding that they require to be the next generation of signal processing innovators. This article describes our work using think-aloud problem-solving sessions to reveal student misconceptions and to design improved conceptual problems for use in active learning classrooms. In addition, we present recently updated Signals and Systems Concept Inventory (SSCI) data demonstrating that courses incorporating active learning pedagogy produce larger gains in students' conceptual understanding of signal processing.

Introduction

Innovation in signal processing requires a deep understanding of core concepts, such as convolution, filtering, and the connection between time and frequency domain representations. Innovation also demands an ability to apply those concepts flexibly to solve novel problems. People who organize ideas within a conceptual framework learn new information quickly and can apply their knowledge in unfamiliar situations [1, p. 17]. According to Montfort et al., conceptual understanding is "more transferable than computational ability" [2]. Those engineering educators define such comprehension as "an understanding of the phenomena underlying a calculation, including the context, purpose, necessary assumptions, and range of reasonable values expected."

While the development of deep conceptual understanding is an intended outcome of science, technology, engineering, and mathematics (STEM) courses, research shows that traditional lecture-based instruction typically does not achieve it [3]. Felder and Brent [4, p. 161] identify three reasons why students can pass a course without attaining a strong conceptual foundation: 1) the instruction does not effectively promote the development of abstract frameworks, 2) the tests don't assess conceptual understanding, and 3) students' misconceptions are too strong to be dislodged. To address these issues, Felder and Brent suggest that a course should include learning exercises designed to

Digital Object Identifier 10.1109/MSP.2021.3060382
Date of current version: 28 April 2021

confront misconceptions and build conceptual understanding, along with exam questions that probe the depth of students' mastery of the concepts.

Active learning [5], [6] is an alternative to traditional lectures that incorporates this approach. Felder and Brent define *active learning* as anything students are asked to do besides observing a lecture and taking notes [4, p. 111]. Active learning incorporates frequent low-stakes assignments and formative assessments, building toward traditional high-stakes exams and summative evaluations [7, p. 139]. Including conceptual problems in both formative and summative assessments encourages students to develop a strong conceptual understanding to complement the computational skills built by more traditional problems. Freeman et al. [3] confirmed that STEM classes that included active learning improved student performance on exams and reduced the number of pupils who failed or withdrew. Active learning also plays an important role reducing performance gaps in STEM classes for students from economically disadvantaged backgrounds [8] and for females in predominantly male courses [9].

A good conceptual question requires minimal computation and provides little or no data to plug into memorized formulas.

This article focuses on improving and assessing students' conceptual understanding of undergraduate signals and systems (S&S), the first course in most signal processing curricula. The core concepts in S&S include linearity and time invariance, convolution, transform representations, filtering, and sampling. The SSCI is one option for summative assessment in S&S courses [10], [11]. This article describes the think-aloud problem-solving interviews that were used to refine and validate the SSCI, along with more recent think-aloud video problems. These sessions led to the development of better open-ended conceptual signal processing problems for formative and summative evaluation.

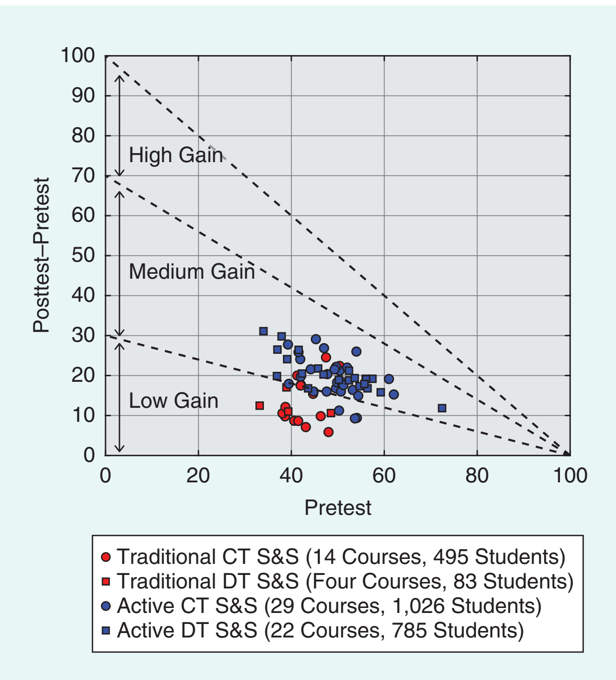

FIGURE 1. SSCI gain data for 2,389 students in 69 courses. Each point represents one course. The abscissa is the SSCI pretest score, and the ordinate is the raw gain, defined as the posttest average minus the pretest average for a course. The low-, medium-, and high-gain regions are those defined by Hake [15]. The average normalized gain for the 18 traditional courses is $\langle g \rangle = 0.23 \pm 0.11$, and for the 51 active learning courses, it is $\langle g \rangle = 0.39 \pm 0.08$.

SSCI: A standardized assessment for S&S

Concept inventories (CIs) are multiple-choice tests designed to assess students' understanding of foundational material in a given field. As described by Epstein in a discussion of the calculus CI [12], "These are tests of the most basic conceptual comprehension of foundations of a subject and not of computation skill." A good conceptual question requires minimal computation and provides little or no data to plug into memorized formulas [10]. The incorrect multiple-choice answers, or distractors, are designed to identify common misconceptions. CIs are often given at the beginning and end of a course to measure students' learning gains [4, p. 164], facilitating comparisons of different pedagogical approaches.

CIs first gained traction in physics with the Force Concept Inventory (FCI) [13] and have since been developed for a variety of disciplines, including engineering [14]. Among these is the SSCI [10], originally developed as part of the National Science Foundation (NSF)-funded Foundation Coalition [14]. Concepts assessed by the 25-question instrument include linearity and time invariance, convolution, frequency representations, and filtering. The first and second authors of this article developed continuous-time (CT) and discrete-time (DT) versions of the SSCI. Spanish translations of both versions are available. The problems discussed in this article are drawn from the CT SSCI.

Hake's influential study of 6,000 students in Newtonian physics courses quantified the positive impact of active learning on pupils' conceptual understanding by using normalized gain, $<g>=(\text{posttest}-\text{pretest})/(100-\text{pretest})$, where *pretest* and *posttest* refer to class average scores on the FCI [15]. Normalized gain can be interpreted as the fraction of the material students learned during a course. Hake found that active learning courses resulted in statistically significantly higher gains than traditional lectures [15]. Using the SSCI, we confirmed that active learning provides similar improvements for S&S classes [6]. Freeman et al.'s [3] meta-analysis of 225 studies of STEM courses, including data from the SSCI, concludes that active learning is the "preferred, empirically validated teaching practice" since it reduces failure rates and boosts exam scores, particularly those for CIs. While the formal SSCI development project concluded in 2009, we continued collecting data of opportunity during the past dozen years. Figure 1 provides updated SSCI gain information from 2,389 students. The average gain across all active learning courses (0.39) is more than a standard deviation above the average across all traditional lecture classes (0.23). This analysis further supports the benefits of active learning and highlights the need for the continued development of better tools for conceptual assessment.

SSCI interview study and other think-aloud exercises

As part of an effort to more deeply understand students' thought processes when solving conceptual problems from the SSCI, in particular, those focused on convolution and filtering, the authors conducted an interview study in which pupils who had recently completed a CT S&S course participated in recorded think-aloud problem-solving sessions. Students were given several questions from the SSCI and asked to talk through their reasoning as they answered each one.

The interviews took place at George Mason University (GMU) and the University of Massachusetts Dartmouth (UMassD) during four semesters, from 2006 through 2008. Forty-one students were interviewed. In a review of best practices for cognitive interviewing, Peterson et al. [16] suggest that a sample size of five to 15 is reasonable as the first step in a validation study, a figure that our study and the other two investigations [17], [18] cited later in this article comfortably exceed. They note that while some problems are not detected until the sample size exceeds 50, the "rate of new problem identification per interview declines, suggesting diminishing returns." While most participants took a CT S&S course the semester before the interview, a few took the class one year earlier. Most were enrolled in a follow-on DT S&S course at the time of the interview. As regional public institutions, GMU and UMassD serve students from a wide range of socioeconomic backgrounds. In the most recent surveys, roughly a quarter of learners at both schools are first-generation college students.

The interviews followed a semistructured protocol; the interviewer asked each participant a common set of questions but was free to probe with follow-up inquiries. The interviewer was an experienced S&S instructor, but no student was interviewed by his or her own teacher. Students received a nominal (US$20) incentive for participating. All questions were from the CT SSCI except for one that was designed specifically for these sessions [19, Fig. 2]. The set of questions participants were asked to solve evolved as the study progressed. The initial cohort of nine students was asked five conceptual questions focusing on frequency-selective filtering and the Fourier transform. We expanded the question pool to include a convolution question for the next three cohorts (nine, six, and eight students). The interviews of the final cohort (nine students) added a new CT SSCI question about convolution.

Initial results of the interview study were published in [19], which describes students' misconceptions about filtering and Fourier transforms, and in [20], which analyzes students' responses to convolution and filtering questions and correlates pupils' SSCI answers with their solutions to similarly themed, open-ended questions on a final exam. One goal of the study was to better understand why learners were selecting (or eliminating) certain answers, in an effort to validate and revise the SSCI. This process sheds light on how broader conceptual assessments (not just CIs) can be designed, implemented, and revised, and it highlights the value of open-ended tasks (including multiple-choice problems with an "explain your reasoning" component) for assessing students' abstract understanding.

Video homework problems provide another avenue for investigating students' comprehension of S&S concepts within the normal flow of a course. Video problems require students to submit a short clip explaining their solution to a homework problem as if they were tutoring a peer. These problems leverage the video recording software widely available on smartphones and tablets. Inspired by a Tweet from Prof. Rhett Allain in a physics education thread, the second author of this article included eight video homework problems in his fall 2020 S&S class. They carried half the points for each assignment, providing incentive for students to do them well. Used this way, video problems provide a tool for instructors to focus students' efforts on the most important topics. Some pupils submitted simple videos narrating their handwritten solution; others produced more elaborate clips. With an enrollment of about 50 students and a 3-min time limit for each video, these problems provided roughly 1,200 additional minutes of students thinking aloud while solving S&S problems, revealing new misconceptions and confirming some previously identified ones. Video problems quickly reveal which pupils actually understand the solution they are presenting.

The interviews and video problems revealed that many students arrived at correct answers despite incorrect and incomplete understanding. The following sections highlight key misconceptions about convolution and filtering that were articulated by students. We present case studies illustrating how we designed improved problems for classes, homework, and quizzes that challenge students to confront misconceptions. Felder and Brent [4] note that these confrontations are a necessary step in the process of developing better conceptual understanding.

Convolution insights

Convolution is a fundamental S&S concept that students often struggle to understand and implement [17], [21]. The think-aloud exercises revealed various misconceptions and highlighted the difficulty of designing convolution questions that probe conceptual understanding instead of memorization and procedural calculus ability. The following discussion describes how the SSCI convolution questions were refined in response to the interviews and presents a new, open-ended problem designed to address misconceptions.

Three cohorts of students in the SSCI interviews answered a single convolution question, Q13 (question numbers refer to the current version of the SSCI, v5), while the last cohort answered two (a modified Q13 and a new Q15). Both problems displayed plots of the impulse response $h(t)$ of a linear time-invariant (LTI) system and the system's input $x(t)$ and asked students to identify the plot of the output $y(t)$. In Q13, $h(t)$ is a unit amplitude rectangular pulse, and $x(t)$ is a unit amplitude square pulse; thus, $y(t)$ is a trapezoid. In the first three cohorts, 21 of 23 students answered Q13 correctly, and the two who answered incorrectly selected the wrong trapezoid, suggesting that they knew to expect that shape but could not accurately determine its parameters. More than half the students who got Q13 correct (12/21) added the starting/ending times of $h(t)$ and $x(t)$ to predict the starting/ending times of the output and select the correct answer, which was the only choice spanning those times. One student referred to this procedure as a *trick*, and others did not justify the approach. The remaining students (nine of 21) who answered Q13 correctly provided some description, with varying levels of detail, of the flip/shift operations required to compute the convolution.

Given the large percentage of students relying solely on predictions of starting/ending points to select an answer, we modified Q13 and designed the new Q15 to probe different aspects of convolution [21]. Those two questions were given to the last cohort (nine students). Analysis of their answers confirmed some of the earlier results and revealed other misconceptions. The modified Q13 posed the original question, but its revised answers were designed to capture whether students knew that they must convolve $h(t)$ and $x(t)$ and that $y(t)$ will be a trapezoid. The correct answer was the only trapezoid among the choices. The majority (seven of nine) answered the revised Q13 correctly. The two who answered incorrectly chose a distractor, $y(t) = h(t) + x(t)$. Interestingly, both of those students mentioned the word *convolution* in their answer; e.g., "Yeah, I mean, I don't know if the correct words are *convolving* or whatever."

The new Q15 defined $h(t)$ and $x(t)$ as unit–amplitude rectangular pulses with lengths of four and two and different starting points. All of Q15's answers had the same starting/ending points, eliminating students' ability to use those cues to select their solution. The correct answer was a trapezoid, and all the distractors were trapezoids or triangles. As expected, Q15 was more difficult, with only four of nine students selecting the correct answer. The most popular distractor (four of five) was the trapezoid with a peak amplitude of one (instead of two). Students justified that choice by saying that the input had an amplitude of one or that $h(t)$ and $x(t)$ were not "scaled by anything." One cited the fact that the output in Q13 had an amplitude of one in justifying that the system did not change the amplitude of a signal, suggesting that the student did not understand what controlled the amplitude of the output. The Q15 responses also illustrated how students extrapolate from the limited examples they have seen. For example, a pupil commented that he had done this type of problem before and that the slope of the trapezoid should be one because "I've almost never seen the slope be anything other than one." While the slope was one for Q15, this is certainly not true in general. Another student expected an answer that started at $t = 0$ because $h(t)$ starts at $t = 0$. This suggests that the student was used to seeing examples where both the input and the impulse response start at $t = 0$.

The SSCI think-aloud interviews provided several insights into students' struggles with convolution. Many learners rely on the "trick" of adding the starting/ending points to determine the output without really understanding why that works. Reflecting on our own classes, we realized that we often use unit–amplitude square and rectangular pulses in our initial set of convolution examples. As the interviews showed, these examples leave some pupils unclear about which signal features control the height and the slope of the output. Similarly, we often use examples where both $x(t)$ and $h(t)$ start at $t = 0$, and this leads some students to assume that all convolutions start at $t = 0$. Also, convolving boxes makes a nice example since it is the easiest integration to do. These interviews motivated us to alter that illustration to make the heights of the boxes equal to A and B so that students explicitly have to multiply A and B together and integrate the result. Then, we give an in-class problem that specifically asks what determined the slope and the maximum height of the output. We have also altered our lecture examples, in-class exercises, and homework problems to include inputs and impulse responses with arbitrary starting and ending points so that students do not assume that convolutions always start at $t = 0$.

Obviously, students need to be able to convolve signals that are not boxes and, most importantly, learn to visualize what the output should look like. Motivated by responses to the Q13/Q15 interviews, we developed the open-ended question shown in "Convolution Question." This problem requires only a simple integration (so it still emphasizes conceptual understanding over computational agility), but it forces learners to think about the shape of the output and what determines the output's maximum height. The problem displays multiple choices for the output and asks students to justify their selection. This question has been used for summative assessment (one of the midterm exams) and subsequently for formative assessment (in-class discussion in small groups).

Filtering insights

Filtering is another fundamental concept in undergraduate S&S. Students need to develop fluency moving between convolving in time and multiplying in frequency. The interviews included three linked CT SSCI problems that probed students' understanding of Fourier transforms and filtering. All 41 pupils were given the filtering problems. We designed those questions to be linked, progressing from frequency-selective filtering (Q6) and frequency representation of narrow-band tones (Q9) to a problem synthesizing both topics (Q25) [10]. This section discusses the questions and students' responses, identifies misconceptions and incomplete models revealed in the think-aloud problem solving, and presents two new open-ended exercises that build on the insights gained about common filtering misconceptions.

Q6 asked students to find the output of an ideal low-pass filter (LPF) when the input is a single-frequency cosine [10]; see "Three-Pulse Filtering Question." (Space constraints preclude reproducing all the SSCI questions from the interviews. We identify previous publications that include the questions for the benefit of interested readers.) The LPF is specified by graphs of the frequency response magnitude and phase, with a gain of three, a cutoff frequency of $\omega = 200$, and a constant phase shift. A large majority (34/41) of the students answered Q6 correctly. However, roughly two-thirds of the students who chose the right answer (22/34) did not explicitly confirm that the signal frequency fell inside the passband before they read the gain and phase shift from the plots provided. In two sessions, the interviewer probed students who chose the correct answer with a follow-up query asking what would happen if the input were $\cos(250t)$ instead of $\cos(50t)$. Both students said the output would be $3\cos(250t)$, failing to recognize that the signal at 250 rad/sample would fall outside the passband and be removed by the filter.

Q9 (numbered Q7 on version 4 of the SSCI [10]) presented students with time-domain plots of two windowed sinusoids, $x_1(t)$ and $x_2(t)$, with $x_2(t)$ being twice the frequency of $x_1(t)$. Given a plot of $|X_1(j\omega)|$, students were asked to identify the plot of $|X_2(j\omega)|$. Most (33/41) answered correctly. Those answering incorrectly were evenly split between students who confused frequency with amplitude and those who properly identified $x_2(t)$ as having a higher frequency than $x_1(t)$ but then confusedly stated that the higher-frequency signal would have its peaks closer together in the frequency domain. Several pupils mentioned that while they knew there was a relationship between time and frequency, they could not remember it. Learners again relied on "tricks" they remembered more than a deep conceptual understanding: "The trick that I learned from one of my classmates was if you see double the cycles, it's just shifted off that much in the

Convolution Question

The following question is designed to prompt students to think about convolution when the integrand is not a constant. It prevents students from relying on starting point and end point tricks to eliminate possible answers and requires a justification.

Question

A linear time-invariant (LTI) system has the impulse response $h(t) = u(t)$ given in Figure S1(a). The input to this system is the signal $x(t)$ in Figure S1(b). Figure S2(a)–(d) displays four signals. Indicate whether each signal could be the output of the LTI system with impulse response $h(t)$ when $x(t)$ is the input. Provide a brief justification for each answer. Answers without justification will receive no credit.

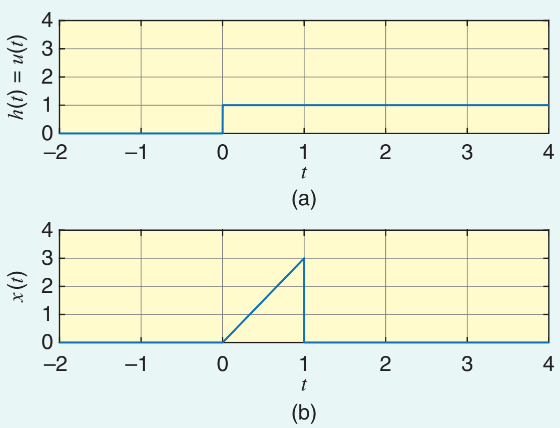

FIGURE S1. (a) The impulse response $h(t) = u(t)$ for the LTI system. (b) The input signal $x(t)$.

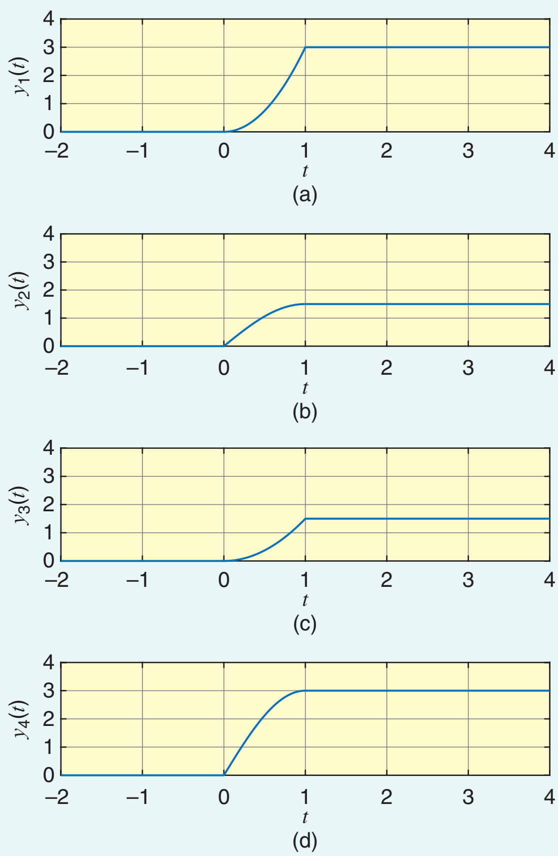

FIGURE S2. (a)–(d) Choices for output of the LTI system when $x(t)$ is the input.

Fourier transform .... I'm not sure what the math is behind that, to tell you the truth."

Q25 [10, Fig. 1] requires students to predict the output signal for an LPF when the input is two narrow-band pulses. The pupils are given graphs of the input time signal $x(t)$ and the Fourier transform magnitude $|X(j\omega)|$ as well as a graph of the filter's frequency response magnitude $|H(j\omega)|$. The LPF's gain is one in the passband. The lower-frequency pulse falls inside the filter passband, but the higher-frequency pulse is eliminated by the filter. As we expected, students found Q25 more difficult than Q6 and Q9: only 28/41 chose the correct answer. Roughly a quarter (eight of 28) of those choosing the right answer demonstrated the same incomplete thought process as in Q6: assuming that the low-frequency pulse would pass through the LPF without ever explicitly multiplying $|X(j\omega)|$ with $|H(j\omega)|$ or even comparing the signal frequencies to the filter cutoff frequency. This incomplete model was especially notable for three students who answered Q25 correctly after they answered Q9 incorrectly. All three made statements identifying the filter as an LPF and stating that the low-frequency pulse would therefore be the output, without any further verification of the filter cutoff frequency.

A large majority of the students who answered Q25 incorrectly (nine of 13) chose the output containing only the high-frequency pulse. Eight of them properly identified the filter as an LPF and explained how to multiply $X(j\omega)H(j\omega)$ in frequency but then linked the low-frequency peak in the filter output $Y(j\omega)$ to the high-frequency pulse in $y(t)$. Many of these learners correctly answered Q9, working forward from time to frequency, but could not apply the same concept in the opposite direction, going from frequency to time.

Students exhibited another misconception in filtering problems when they "masked" the input spectrum with the frequency response rather than multiplying the two. Learners confused about filters in this way treat the passband like a "mask" or cookie cutter that trims away any of the input Fourier transform falling outside the filter passband in frequency or amplitude. One justified choosing the Q6 distractor with the gain of one instead of three by comparing the frequency of the input cosine to the cutoff frequency of the LPF, also explicitly noting that the cosine amplitude of one was less than the filter passband gain of three. The same misconception appeared in roughly 15% of the answers submitted for a fall 2020 video problem about filtering. Students reasoned that the filter would only pass those parts of the input spectrum "inside" the passband in both frequency and amplitude. This misconception demonstrates the value of think-aloud problems. The masking misconception was not one we identified prior to drafting the SSCI.

Reflecting on the insights from the interviews led us to design new problems for in-class exercises, homework, and exams. Those problems were often open-ended to enable us to

Three-Pulse Filtering Question

The following open-ended synthesis question is designed to prompt students to explicitly consider filter cutoffs. The question has a range of possible solutions, better reflecting the engineering design problems students will face as practicing professionals.

Question

Let $x[n]$ be a discrete-time (DT) signal with three pulses. Figure S3 displays $x[n]$ using Matlab's `plot` command. The signal is actually DT but is presented using `plot` instead of stem to reduce clutter. The DT Fourier transform magnitude for $x[n]$ appears in Figure S4. When the signal $x[n]$ is the input to an LTI filter, the output is the signal $y[n]$ in Figure S5 (again, using `plot` instead of `stem`).

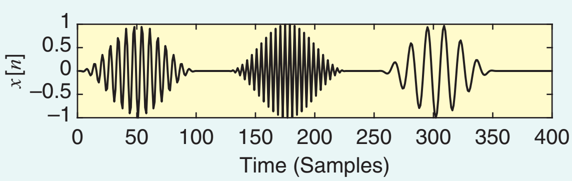

FIGURE S3. The DT signal $x[n]$.

Sketch a frequency response magnitude $|H(e^{j\omega})|$ consistent with the information given. There may be more than one correct answer. Your sketch should cover the frequencies $0 \le \omega \le \pi$ and must label all important frequencies and amplitudes. Write a short explanation (three sentences or fewer) to receive full credit.

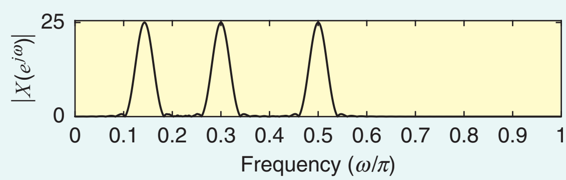

FIGURE S4. The DT Fourier transform magnitude for $x[n]$.

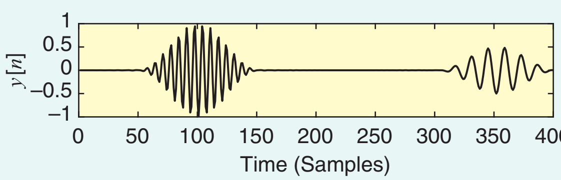

FIGURE S5. The output signal $y[n]$ when the signal $x[n]$ is the input to an LTI filter.

observe unanticipated misconceptions. To encourage students to compare filter cutoff frequencies with input signal frequencies, we designed new filtering questions similar to Q25 but using three narrow-band input pulses and asking learners to draw the frequency response of a filter that could produce the given output; see "Three-Pulse Filtering Question." Students cannot deploy simplistic strategies, as many did on Q25, but must explicitly consider the middle-frequency pulse. This problem also pushes pupils to think more deeply about the role of gain, in contrast with ideal frequency-selective filters, where the gain is always one or zero. In addition to the version shown here, we also asked "three-pulse" filtering questions in which the output contained one, three, or zero pulses. Students often find a zero output uncomfortable, as few textbooks and class examples portray this case. This discomfort opens the door to a deeper discussion of the role of filters in practical systems. Another variation on the filtering theme of Q25 graphed $y[n]$, $|Y(e^{j\omega})|$, and $|H(e^{j\omega})|$ and asked students which of six choices for the input $x[n]$ were consistent with the information given. In all these problems, learners were required to justify their answers with two or three sentences to receive full credit.

As instructors, we have a responsibility to prepare students to explain and defend technical choices to their peers

To address the masking misconception, filtering problems must include amplitudes in the input transform $X(j\omega)$ that are greater than the gain of the frequency response $H(j\omega)$ at some frequencies. "Masking Question" illustrates one such problem written in response to the interviews. The difference between multiplying and masking will be very clear in how a student's sketch of $Y(j\omega)$ treats the triangle in $X(j\omega)$ in the region $-50 < \omega < 50$. Pupils filtering correctly will sketch the output in this frequency band to be a triangle with a peak of six. Students incorrectly masking will remove any of the triangle above two, leaving a large, flat plateau. Similarly, the output transform $Y(j\omega)$ will exhibit dramatic differences between the correct and incorrect approaches in other frequency bands.

Linking our findings to prior studies

Two other studies interviewed S&S students in think-aloud problem solving similar to the SSCI interviews [17], [18]. The problems students solved in those investigations combined conceptual and procedural knowledge and did not satisfy the definition of a concept question, given in [10], that one can identify the correct answer without requiring paper and a pencil. Nasr et al. [17] interviewed 24 students in the linear systems module of an aeronautical engineering course. Each pupil solved one of three problems about superposition and convolution in LTI systems. Jia et al. [18] asked 24 students to solve four problems about the trigonometric form of the CT Fourier series in a 1-h interview. As described in the following paragraphs, several of our findings match those identified by Jia and Nasr.

Even though Jia et al. [18] focused on the Fourier series, they observed underlying challenges for students connecting time and frequency representations that were similar to those we found in the filtering problem. More broadly, they also identified students who relied on rote procedural techniques and tricks that left them adrift when a problem required more flexibility in exploiting the information given. Last, they found that students suffered from a "disconnect" when translating between plots and equations, echoing comments from our interviews. Nasr et al. [17] found that roughly half the students struggled with each of the interview problems covering convolution. They observed a version of the masking misconception in the context of multiplying signals within the convolution integral. Students solving a graphical "flip-and-slide" convolution problem stated that the integrand $x(\tau)h(t-\tau)$ would be zero for the region in time where the input was positive and that the impulse response was negative because when one signal was above the axis and the other was below the axis, they did not "overlap." We have

Masking Question

The following question is intended to unearth students' masking misconceptions in filtering. If learners treat the frequency domain operation as masking rather than multiplication, the peak of each triangle will be clipped in the output Fourier transform.

Question

A continuous-time LTI system has an input $x(t)$ and an output $y(t)$. The frequency response $H(j\omega)$ of the system is plotted in Figure S6. The Fourier transform $X(j\omega)$ of the input is shown in Figure S7. Determine and sketch the Fourier transform $Y(j\omega)$ of the system output.

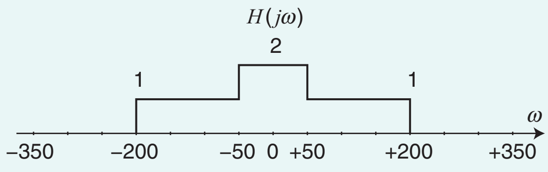

FIGURE S6. The system's frequency response $H(j\omega)$.

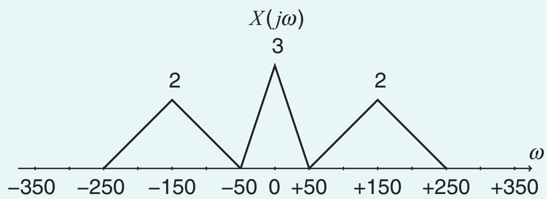

FIGURE S7. The input's Fourier transform $X(j\omega)$.

Phil and Connie

This story problem challenges students to confront misconceptions about causality. The second author of this article gratefully acknowledges Abelson and Sussman's story problems in their Structure and Interpretation of Computer Programs course for inspiring him to create similar signals and systems questions.

Problem

Two classmates, Phil Turing and Connie Volution, are arguing about whether the system $y[n]=x[n^2]$ is causal. You must settle the argument.

Phil: "The system is causal. If the input $x[n]=\delta[n]$, the output $y[n]=h[n]$. Substituting $\Delta[n]$ and $h[n]$ into the system equation gives $h[n]=\delta[n^2]$. From this, we see that $h[n]$ is one when $n^2=0$, which is $n=0$, and otherwise, $h[n]=0$ for all other values of n. Since $h[n]$ is zero for negative time, the system must be causal."

Connie: "The system is not causal. Consider the output when $n=2$. Then, $y[2]=x[2^2]=x[4]$. The output at this value of n depends on the input $x[n]$ at a future time sample, so the system must not be causal."

Which student is right, Phil or Connie? What mistake did the incorrect student make?

also observed that many students struggle with correctly interpreting the point-by-point multiplication of two signals presented graphically.

Conclusions

This article summarized what we learned about students' S&S misconceptions from think-aloud interviews and video homework problems. We presented examples illustrating how we changed our teaching and assessments to challenge learners to develop deep conceptual understanding. In this section, we highlight two key insights from our analysis and encourage fellow instructors to experiment with think-aloud problems in their own courses.

The first insight is that many students seize on "tricks" as shortcuts to avoid developing deep conceptual understanding. Students repeatedly referred to tricks to describe procedures to solve convolution and filtering problems. Rather than recognizing these procedures as efficient approaches to common tasks, derived from S&S theory, learners treated them as magic spells to be memorized. Instructors may inadvertently encourage magical thinking when they use informal language to define and derive the procedure. How do we address this issue? Ambrose et al. [7, p. 3] remind us that "learning is not something done *to* students, but rather something students themselves do." The key to getting students to stop viewing procedures as "magic" is to have learners derive and test methods themselves by working through a series of linked problems and then to ask them to reflect on what they learned from their answers.

Story problems are another technique to force students to focus on the assumptions underlying tricks. These exercises require pupils to adjudicate disagreements between fictional peers, some of whom articulate different misconceptions. "Phil and Connie" shows a story problem we developed to help students understand causality in LTI systems. Story problems fit neatly within Felder and Brent's framework for addressing misconceptions, as they require students to commit to one view point and to confront a contradiction that ensues [4]. Story problems bring the additional benefit of challenging learners at a higher level of Bloom's taxonomy [22], calling for the evaluation and analysis of the concepts tested, not just application and understanding.

The second insight is that the examples we choose to illustrate S&S concepts have major consequences for what students learn. As instructors, we simplify examples to reduce the cognitive load when first presenting ideas. To teach convolution, we often focus so much on examples with rectangular pulses with unit amplitude that students conclude that the slope of the result has to be one since they haven't seen anything else. When we design filtering problems, we always have at least one signal removed by the filter, enabling students to develop bad habits, such as failing to compare signal frequencies to passband edges. As instructors, we need to assemble comprehensive sets of exemplar problems to enable students to develop complete and accurate mental models. We also need to be mindful of how our problem sets might inadvertently present patterns we do not intend and cause pupils to overgeneralize. Ideally, we want students to be able to generalize their knowledge beyond the exemplars seen in a course. Active learning and think-aloud problems enable instructors to observe while students construct mental models and to prevent learners from "overfitting" to unimportant details.

We encourage our fellow S&S instructors to experiment with more open-ended conceptual problems and to include think-aloud exercises in their classes. First, asking students to explain their thinking orally presents an opportunity for instructors to learn more about learners' reasoning, providing valuable feedback to correct misconceptions before they are entrenched, and improve future versions of a course. Second, these oral explanations challenge students to think about what they understand. Explaining a procedure to someone else helps pupils comprehend those procedures more deeply rather than simply going through the steps silently on their own. Orally explaining a solution to a problem (or sequence of problems) can challenge students to make connections between concepts and procedures and hence develop a more complete grasp of complex ideas. Finally, students' future careers will revolve around oral discussions during design reviews and team meetings. As instructors, we have a responsibility to prepare them to explain and defend technical choices to their peers. Students combining deep conceptual understanding with convincing verbal explanations will be primed to be the next generation of signal processing innovators.

Authors

Kathleen E. Wage (k.e.wage@ieee.org) received her Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology/Woods Hole Oceanographic Institution Joint Program. She is an associate professor of electrical and computer engineering at George Mason University, Fairfax, Virginia, 22030, USA. She received the 2016 IEEE/Hewlett-Packard Harriett B. Rigas Award and the 2008 Mac Van Valkenburg Early Career Teaching Award and is a coauthor of the Signals and Systems Concept Inventory. Her research interests include array processing, random matrix theory, underwater acoustics, and signal processing education. She is a Senior Member of IEEE.

John R. Buck (johnbuck@ieee.org) received his Ph.D. degree from the Massachusetts Institute of Technology/Woods Hole Oceanographic Institution Joint Program. He is a chancellor professor of electrical and computer engineering at the University of Massachusetts Dartmouth, North Dartmouth, Massachusetts, 02747, USA. He is a fellow of the Acoustical Society of America, the 2005 recipient of the Mac Van Valkenburg Early Career Teaching Award, and a coauthor of the Signals and Systems Concept Inventory. His research interests include array signal processing, animal bioacoustics, and engineering pedagogy. He is a Senior Member of IEEE.

Jill K. Nelson (jnelson@gmu.edu) received her Ph.D. degree in electrical engineering from the University of Illinois at Urbana–Champaign. She is an associate professor of electrical and computer engineering at George Mason University, Fairfax, Virginia, 22030, USA. She is a 2010 recipient of the National Science Foundation CAREER Award and the 2014 recipient of the Mac Van Valkenburg Early Career Teaching Award. She is also a member at large of the IEEE Education Society Board of Governors. Her research interests include detection, estimation, and machine intelligence for sonar applications, as well faculty teaching development through communities of practice. She is a Member of IEEE.

Margret A. Hjalmarson (mhjalmar@gmu.edu) received her Ph.D. degree in curriculum and instruction, with a concentration in mathematics education, from Purdue University. She is a professor in the School of Education, George Mason University, Fairfax, Virginia, 22030, USA. Her research interests include engineering education and mathematics education.

Acknowledgments

This material is based upon work supported by the NSF, under grants DUE-0512686, DUE-0512430, and EEC-9802942. This article was written while Margret A. Hjalmarson and Jill K. Nelson served as NSF program officers. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF or Jean-Baptiste Joseph Fourier.

References

[1] J. D. Bransford, A. L. Brown, and R. R. Cocking, Eds., *How People Learn: Brain, Mind, Experience, and School, Expanded Edition*. Washington, D.C.: National Academy Press, 2000.

[2] D. Montfort, S. Brown, and D. Pollock, "An investigation of students' conceptual understanding in related sophomore to graduate-level engineering and mechanics courses," *J. Eng. Educ.*, vol. 98, no. 2, pp. 111–129, 2009. doi: 10.1002/j.2168-9830.2009.tb01011.x.

[3] S. Freeman, S. L. Eddy, M. McDonough, M. K. Smith, N. Okoroafor, H. Jordt, and M. P. Wenderoth, "Active learning increases student performance in science, engineering, and mathematics," *Proc. Nat. Acad. Sci. (PNAS)*, vol. 111, no. 23, pp. 8410–8415, 2014. doi: 10.1073/pnas.1319030111.

[4] R. M. Felder and R. Brent, *Teaching and Learning STEM: A Practical Guide*, San Francisco: Jossey-Bass, 2016.

[5] J. R. Buck, K. E. Wage, and J. K. Nelson, "Designing active learning environments," *Acoust. Today*, vol. 12, no. 2, pp. 12–20, 2016.

[6] J. R. Buck and K. E. Wage, "Active and cooperative learning in signal processing courses," *IEEE Signal Process. Mag.*, vol. 22, no. 2, pp. 76–81, Mar. 2005. doi: 10.1109/MSP.2005.1406489.

[7] S. A. Ambrose, M. W. Bridges, M. DiPietro, M. C. Lovett, and M. K. Norman, *How Learning Works: Seven Research-based Principles for Smart Teaching.* San Francisco: Jossey-Bass, 2010.

[8] D. C. Haak, J. HilleRisLambers, E. Pitre, and S. Freeman, "Increased structure and active learning reduce the achievement gap in introductory biology," *Science*, vol. 332, no. 6034, pp. 1213–1216, 2011. doi: 10.1126/science.1204820.

[9] M. Lorenzo, C. H. Crouch, and E. Mazur, "Reducing the gender gap in the physics classroom," *Amer. J. Phys.*, vol. 74, no. 2, pp. 118–122, 2006. doi: 10.1119/1.2162549.

[10] K. E. Wage, J. R. Buck, C. H. G. Wright, and T. B. Welch, "The signals and systems concept inventory," *IEEE Trans. Educ.*, vol. 48, no. 3, pp. 448–461, 2005. doi: 10.1109/TE.2005.849746.

[11] http://www.signals-and-systems.org

[12] J. Epstein, "The Calculus concept inventory – Measurement of the effect of teaching methodology in mathematics," *Notices Amer. Math. Soc.*, vol. 60, no. 8, pp. 1018–1026, 2013. doi: 10.1090/noti1033.

[13] D. Hestenes, M. Wells, and G. Swackhamer, "Force Concept Inventory," *Phys. Teacher*, vol. 30, no. 3, pp. 141–158, 1992. doi: 10.1119/1.2343497.

[14] D. L. Evans, G. Gray, S. Krause, J. Martin, C. Midkiff, B. Notaros, M. Pavelich, D. Rancour et al., "Progress on concept inventory assessment tools," in *Proc. 33rd ASEE/IEEE Front. Educ. Conf.*, Nov. 2003, pp. T4G1–T4G8.

[15] R. R. Hake, "Interactive-engagement vs. traditional methods: A six-thousand-student survey of mechanics test data for introductory physics courses," *Amer. J. Phys.*, vol. 66, no. 1, pp. 64–74, Jan. 1998. doi: 10.1119/1.18809.

[16] C. H. Peterson, N. A. Peterson, and K. G. Powell, "Cognitive interviewing for item development: Validity evidence based on content and response processes," *Meas. Eval. Couns. Dev.*, vol. 50, no. 4, pp. 217–223, Oct. 2017. doi: 10.1080/07481756.2017.1339564.

[17] R. Nasr, S. R. Hall, P. Garik, "Student misconceptions in signals and systems and their origins – Part II," in *Proc. 35th ASEE/IEEE Front. Educ. Conf.*, Oct. 2005, pp. T4E-26-31.

[18] C. Jia, A. G. Bennett, D.-H. Nguyen, N. S. Rebello, and S. Warren, "Teaching-learning interviews to understand difficulties with Fourier series concepts," in *Proc. ASEE Annu. Conf.*, June 2011, pp. 22.1409.1–22.1409.16.

[19] K. E. Wage, J. R. Buck, and M. A. Hjalmarson, "Analyzing misconceptions using the signals and systems concept inventory and student interviews," in *Proc. IEEE DSP/SPE Workshop*, 2006, pp. 123–128.

[20] J. R. Buck, K. E. Wage, M. A. Hjalmarson, and J. K. Nelson, "Comparing student understanding of signals and systems using a concept inventory, a traditional exam and interviews," in *Proc. 37th ASEE/IEEE Front. Educ. Conf.*, Oct. 2007, S1G1–S1G6.

[21] K. E. Wage, J. R. Buck, J. K. Nelson, and M. A. Hjalmarson, "Signals and systems assessment: Comparison of responses to multiple choice conceptual questions and open-ended final exam problems," in *Proc. IEEE SPE Workshop*, 2011, pp. 198–203.

[22] B. S. Bloom Ed., *Taxonomy of Educational Objectives: Handbook 1: Cognitive domain*, New York: McKay, 1956.

SP

Jorge Torres Gómez, Antonio Rodríguez-Hidalgo,
Yannelys Virginia Jerez Naranjo, and
Carmen Peláez-Moreno

Teaching Differently

The digital signal processing of multimedia content through the use of liberal arts

©SHUTTERSTOCK.COM/LUCKYSTEP

Generally, the curriculum design for undergraduate students enrolled in digital signal processing (DSP)-related engineering programs covers hard topics from specific disciplines, namely, mathematics, digital electronics, or programming. Typically, these topics are very demanding from the point of view of both students and teachers due to the inherent complexity of the mathematical formulations. However, improvements to the effectiveness of teaching can be achieved through a multisensorial approach supported by the liberal arts. By including the development of art and literacy skills in the curriculum design, the fundamentals of DSP topics may be taught from a qualitative perspective, compared to the solely analytical standpoint taken by traditional curricula. We postulate that this approach increases both the comprehension and memorization of abstract concepts by stimulating students' creativity and curiosity. In this article, we elaborate upon a methodology that incorporates liberal arts concepts into the teaching of signal processing techniques. We also illustrate the application of this methodology through specific classroom activities related to the digital processing of multimedia contents in undergraduate academic programmes. With this proposal, we also aim to lessen the perceived difficulty of the topic, stimulate critical thinking, and establish a framework within which nonengineering departments may contribute to the teaching of engineering subjects.

Introduction

Engineering courses related to signal processing rely on strong mathematical methods, such as linear-time-invariant systems, complex analysis, or transform theory, among others [1]. These are topics that are often very difficult to grasp. They are also difficult to teach; the bottleneck is, most of the times, the complexity of the abstract mathematics underpinning the core concepts as perceived by the students. The development of attractive and effective comprehension methods to enlighten the fundamentals plays a major role in facing the natural intellectual barriers in this highly specialized field. Thus, the application of new methodologies used to teach core concepts, which

Digital Object Identifier 10.1109/MSP.2021.3053218
Date of current version: 28 April 2021

 1053-5888/21©2021IEEE

take a multisensorial approach through the use of liberal arts, may trigger self-motivation and critical-thinking abilities, allowing students to discover the meaning and implications of their analytical formulations.

Liberal arts education is commonly related to the discipline of social sciences, while mathematics and physics are typically addressed in pure sciences and electrical engineering programs. Consequently, a historical divide has been established in schools that separate arts from science, forming to two markedly different cultures: literacy-based intellectuals and natural scientists [2]. Mixing the schools, however, brings a variety of benefits worth considering. A better education, achieved by bringing the two branches together, will empower students, offering a more complete understanding of the role played by technology and its impact on society (this subject is largely ignored by current academic programs [3]). According to [4], educational skills related to the arts of reading, writing, and communication as well as the appreciation of fine arts could contribute to an improvement in critical thinking and creativity, leading to better proficiency for engineers in their respective fields.

Blending inductive and deductive approaches leads to the formation of powerful skills that can be used to deal with the challenging problems facing society today.

On the one hand, it is commonly accepted that the inductive experience of appreciating and producing a work of art stimulates the production, discovery, and invention of novel technical solutions. However, deduction processes (which are commonly addressed in traditional courses) are more concerned with the particular application of a general theory in the context of engineering applications. Blending inductive and deductive approaches leads to the formation of powerful skills that can be used to deal with the challenging problems facing society today. In this regard, the application of art concepts into the academic engineering program may lay the foundations for new solutions, stimulating both inductive and deductive abilities [4]–[6].

In addition, the transition from science, technology, engineering, and mathematics (STEM) to science, technology, engineering, arts, and mathematics [7], [8] has been advocated by several authors as a way to break down the distinction between disciplines traditionally seen as "creative" (like the arts or music) and STEM disciplines, which are usually seen as more rigid or strictly logical [9]–[12]. However, most of this research studies school education and rarely includes university-level classes. Moreover, this destruction of artificial boundaries could be a way to reduce the persistent gender gap found in STEM studies [13], [14], improving female enrollment in STEM courses and, following this, in physics, engineering, and computer-science-related careers [15], [16].

Today, most engineering programs do not effectively integrate both cultures. Mixed academic programs provide courses that cover both technology and the appreciation of art. Some of these approaches allow engineering students to choose subjects offered by different departments, thereby promoting an interdisciplinary dialogue between arts and engineering [3]. For instance, some programs motivate students to take part in cultural activities, such as museum visits or concerts [17]; they even offer the possibility to participate in academic projects that are coached by professors from both the social sciences and STEM departments [18]. Other programs integrate arts and animation when teaching computer programming to increase students' motivation to learn [19]. In addition, through the use of emoticoding, the integration of algebra, geometry, music, and 3D art has been shown to improve student learning as well [20]. Another example involves web designers, who cleverly combine technology, arts, and psychology to attract the attention of potential customers [21]. Finally, in [22], the authors present abstract mathematical objects and concepts that have been turned into visual art using mathematical animations, thereby offering students a more tangible experience that can potentially improve and enrich their understanding.

Including arts and their corresponding inductive processes into engineering programs is a challenging endeavor. The aforementioned approaches illustrate useful points of contact; however, the effect of a tighter integration of the arts with the teaching of core concepts is still unexplored. Our current research is part of an ongoing teaching innovation project that improves the understanding of engineering concepts using arts as the motor core of a more complex process. It is hoped that this will eventually lead to greater student engagement with their own learning development. When students are encouraged to interact with artistic production goals using engineering concepts, their internalization of key engineering topics is stimulated. Ambitiously, we intend to reinforce the understanding of key concepts by promoting the appreciation and creation of small pieces of art during the teaching of engineering curricula.

In this article, we elaborate on a methodology that integrates art and engineering to more effectively teach core engineering concepts. We illustrate the use of the proposed methodology through specific classroom activities involving the use of signal processing techniques applied to multimedia coding. Using this approach, different engineering topics will be taught from a subjective standpoint. The tools used will include an elaboration on narrative excerpts and their auditory representation by means of well-known musical scales and the processing of images by appreciating their composition. The students are prompted to write their own thoughts and opinions about the topic at hand, and their textual contribution is then used as the media upon which DSP concepts are later implemented. In this way, students will have a multisensorial experience when learning the core concepts in their field. The meaning behind their analytical definitions and implications will be exposed by the analyses and production of art created by writing, listening, and visualization.

We have chosen to employ this methodology in a telecommunications engineering program in which DSP techniques play a major role. By doing so, we seek to provide proof that this integration can be effectively implemented. We also aim to offer inspiration and motivation for colleagues teaching similar concepts, but also totally different, to use nonengineering disciplines to educate better and more well-rounded engineering professionals. Some hints and suggestions regarding the further application of the methodology to other DSP-related subjects are also discussed.

Methodology

As discussed in the previous section, the introduction of art and critical thinking into an engineering curriculum is beneficial to the development of more creative and socially aware engineers and could also promote vocations in STEM careers for women [15], [16]. This is, however, a challenging enterprise, mainly because the materials and resources need to be developed ad hoc. Some general principles, applicable to any subject, need to be observed during this process. In this regard, we have identified the following steps:

1) A small set of key or core concepts, underpinning the subject being taught, need to be identified. These should be the main concepts that students would need to retrieve from their memory in their future professional lives when facing new unforeseeable challenges. Therefore, the instructors' goal is to find ways of making these concepts more suggestive, thereby reinforcing their persistence in memory.
2) Original and suggestive connections between the previous key or core concepts and arts must be found. This is a creative process that considers the connections' motivational ability, where the elements of surprise and originality should be sought. It is a critical step that relies on the well-established benefits of surprise, i.e., the emotional response to outcomes that do not match our expectations [23], to enhance learning [24]. Neuroscientific literature suggests that the prediction errors caused by this mismatch play a universal role in driving learning throughout the human brain [25], increasing attention that, in turn, leads to more effective memorization [26].
3) A set of activities to explore the previous connections needs to be designed. A requirement of these activities is that they promote critical thinking and creativity; that is, the application of mechanistic rules and closed solutions should be avoided. A proper balance of analysis and synthesis procedures should be sought, promoting both deductive and inductive abilities in the students. In addition, these activities should be partially aligned with the multimedia learning theory proposed by the authors in [27] because our students will be learning concepts that include both verbal and written content and more visually and auditory rewarding ideas. This will promote a more in-depth understanding.
4) An assessment procedure tailored to evaluate the degree to which the core concepts are understood needs to be put in place.

In the following section, we illustrate the application of this methodology to enhance the learning of source coding theory in the field of communication systems.

A historical divide has been established in schools that separate arts from science, forming two markedly different cultures.

Description of the academic course

We set out to apply the steps defined in the "Methodology" section to "Multimedia Information Coding for Communications," a compulsory subject in the third year of a four-year bachelor's degree in mobile and space communications engineering, taught at Universidad Carlos III de Madrid. The main goal of the course is to provide an understanding of the coding and compression techniques used to process digital multimedia content such as speech, audio, image, and video to reduce their storage requirements, equivalently, their transmission bit rate.

The original course comprises 14 lectures, with five seminars to develop problem-solving skills and eight 2-h labs. The exercises we describe in the "A Proposal of Activities Regarding the Engineering Concepts and the Arts" section have been implemented during three of the lecture's time slots. The academic program addresses common coding techniques that are based on the statistics or entropy of the source (Huffman, arithmetic, Golomb–Rice, and Lampel–Ziv) and on the perception of audio, images, and video (perceptual coding) to represent and transmit data. In addition, the evolution of the most relevant standards is presented, promoting discussion of the coding performance of current applications; these standards include MP3 for audio compression, JPEG for image, and H.26x or the MPEG family for video (see, for example, [28] and [29]). In addition, during the lab sessions, students develop hands-on skills by programming coding algorithms using MATLAB.

The program surveys the following six topics:

1) *Fundamentals of the digitalization of multimedia information*: This topic provides a quick review of the fundamentals of analog-to-digital conversion; covers the representation by sequences, matrices, and time-varying matrices of audio, image, and video, respectively; and gives an overview of the fundamentals of compressing information. The students are already familiar with these concepts, so this is an introductory topic.
2) *Speech coding*: This topic establishes the mechanisms of human speech production as a way to introduce the basis of coding. Vocoder, hybrid, and waveform technologies are presented (together with their efficiency tradeoffs) for speech transmission over communication channels, such as those of mobile communication networks.
3) *Audio coding*: This topic covers a broader representation of sound in comparison to speech. Additional codes and a first survey of the related standards are taught. The

concept of perceptual coding is also explained, based on a model of human hearing and associated psychoacoustical phenomena.

4) *Image coding*: This topic includes the analysis of 2D data in contrast with the previous two sections, where audio and speech were considered. 2D transforms are presented as a way to apply coding in a different domain to reduce storage and transmission demands. The mixed solution of transformation and entropy coding is presented, together with the evolution of standards.
5) *Video coding*: This topic introduces how redundancy is reduced in video through the efficient processing of the evolution of its frames to account for motion in the sequence of images. Video formats, general coding strategies, and the related standards are illustrated in this section.
6) *Transmission of multimedia information*: This final topic covers the requirements for universal access to multimedia. Technologies and methods are presented by means of illustrative approaches, standards and network protocols.

Topics 2–5 are accompanied by lab assignments to implement very simplified examples of each of the coder types where reference decoder implementations are provided. The student's goal is to produce coder implementations whose outputs can be properly decoded using the supplied reference decoders, as would be the situation in an organization that develops coding standards. Therefore, all of these topics deal with the digital representation and processing of multimedia information. Implicitly, students need to implement basic DSP techniques, such as filtering, downsampling, upsampling, and transformations, to handle the proper functioning of these coders.

Generally speaking, multimedia content is a data source with specific characteristics or patterns that may be exploited for compression purposes. For instance, the speech recording of a tenor voice has major differences when compared to the recording of a bass one. The fundamental frequency of the first is higher than that of the second. A similar distinction can be observed in audio signals because classical music presents different timbres and rhythms compared to jazz or soul. Images, on the other hand, exhibit huge differences in color and patterns depending of the age or art movement of their creation: paintings from the Baroque period contrast strongly with Cubist artwork, for example.

Coding techniques may exploit the peculiarities of media sources to give a better performance in terms of the data storage—or the equivalent rate of transmission—needed for a given quality of reproduction. This connection between the nature of the media source and the coding performance—either in terms of the compression rate or the quality (or distortion) of the outcome—is mediated by the concept of redundancy or by the amount of information as measured by the entropy of the source. As the importance of these ideas traverses all of the topics on this subject, they will be chosen as the core concepts to be identified in the first step of the methodology defined in the "Methodology" section. The next section is devoted to illustrating the teaching of these core concepts through the design of specific hands-on exercises that relay on the analysis and synthesis of small pieces of art.

A proposal of activities regarding the engineering concepts and the arts

An appreciation of art may produce a better understanding of coding concepts and applications related to the digital processing of multimedia content, provided that original connections involving elements of surprise are included, as explained in the "Methodology" section. In the "Description of the Academic Course" section, we identified entropy and redundancy as core concepts (step 1). We now need to find suggestive connections (step 2) between these concepts and a certain modality of art and develop open and creative activities that will make these connections apparent (step 3).

To this end, the proposed exercises are divided into two separate parts designed to foster both deductive and inductive skills: first, illustrating the differences in coding performance through the appreciation of representative pieces of art from a qualitative standpoint (analysis), and second, challenging students to produce art based on the concepts and fundamentals of coding, thereby becoming part of an active process (synthesis). This organization resonates with several extremely important DSP concepts, such as dualities in linear transformations of sequences, and reinforces the complementarity of coding and decoding stages, which is essential in this subject. The analysis and synthesis assignments will involve the implementation of DSP techniques to process text, sounds, and images. The following exercises describe these two directions in the topics of speech, audio, and image coding.

Exercise 1: Compressing the concept of entropy

This exercise is related to the concept of entropy. The students will produce a brief essay, preferably related to their understanding of entropy. The whole process of writing, in and of itself, will make the students focus on the task at hand.

Once their assignments are written, the students are asked to read and record them into a raw audio file. Note that, semantically speaking, the informational content of both text and audio is the same, albeit in totally different modalities. Note here also that the vagueness of the definition of information, a concept new to the students at this stage, is sought to trigger discussions that they must solve if they are to attain a solution.

In particular, the student should derive the following (analysis phase): 1) the entropy of the written script, 2) the entropy of the recorded speech, and 3) a comparison and a discussion based on the results of tasks 1 and 2. In the second stage of this exercise, students must produce a new text reflecting the same content but with fewer words (synthesis phase). Note that this is a procedure of compression itself as students remove what they think is irrelevant or less informational in the texts that they have produced. The students must complete tasks 1–3 again.

In addition to understanding the concept of entropy, during the discussion, the students are expected to express their opinions about which of these two formats (the first or the second definition of entropy they have produced) seems to be more useful and comprehensible. This whole procedure is designed to help them consider the usefulness of this process; that is, whether or not the concept is understandable after removing some of the irrelevant text, whether or not a person unrelated to engineering is capable of understanding it before and after the removal of redundant information, and to what extent uninformative data can be removed before the concept becomes unintelligible.

Additionally, students will develop DSP-related technical skills when implementing solutions to processing pieces of text. That is, obtaining histograms (when computing the entropy), recognizing patterns (when identifying the appearance of letters in the written text), and estimating their corresponding probabilities, and so on.

A requirement of these activities is that they promote critical thinking and creativity; that is, the application of mechanistic rules and closed solutions should be avoided.

Exercise 2: Listening to redundancy

This exercise proposes producing (nonspeech) sounds from a given written script. Globally, this could be referred to as *sonification*, with applications not only in text-to-sound processes but also to images-to-sound and vice versa. Here we focus on the production of sound from a written script to gain a subjective measure of redundancy, another core concept of coding theory.

Redundancy concerns the existence of multiple replicas of a given attribute. A good coder will try to reduce the redundancy of the coded sequence to produce a more compact representation, benefiting from the fact that a redundant attribute (that is, an attribute with a high a priori probability) will be easier to predict.

In terms of sound, this may be reflected in multiple ways; for example, by assigning a chord to each of the symbols (i.e., each letter) used to code a text. If a particular letter is highly dominant or redundant in a text, this will be noticeable; it may even be interpreted as the tonic chord. If the given text presents too many different attributes that are almost identically distributed, the music produced would sound much more chaotic. However, if there is a high number of symbols, the sonification will be difficult to perceive and achieve.

Therefore, in this exercise, we resort to a more nuanced sonification method based on the perception of the different timbres of a sound; in other words, the quality that allows us to distinguish between different musical instruments (related to the shape of the envelope of the spectra of the produced notes). A text with highly redundant symbols will sound as a different instrument than one where the symbols are equally redundant. Moreover, the reduction of the redundancy produced by an effective coder (in comparison to a less-effective one) could also be perceived as a different instrument.

Based on this interpretation, the analysis phase of this exercise contains the following steps:

1) to produce a written script by explaining some of the learned psychoacoustic principles. Specifically, the script must provide definitions for the following aspects (of no fewer than 15 words per item): psychoacoustic principle, the absolute threshold of hearing, time masking, and frequency masking
2) to encode the produced text and obtain a corresponding histogram of the produced symbols for two different coders: fixed length and Huffman
3) to apply sonification techniques to produce sounds from the two obtained histograms
4) based on the produced sounds to derive concluding remarks regarding the capacity to reduce redundancy through the use of two coding techniques.

To illustrate a possible sonification strategy for item 4 we start with a histogram representation of $L_i \cdot p_i$, where L_i is the length of the code obtained for symbol i and p_i represents the a priori probability, as depicted in Figure 1(a) and (c) for the fixed- and variable-length codes, respectively, used to code the letters of the alphabet. The histogram bars will be the attributes used to convert text into sound. The students must implement the following steps:

1) Reorder the histogram bars in decreasing order as depicted in Figure 1(b).
2) Interpret them as a Fourier series in which the produced sound will be the linear combination of pure tones that are harmonically related to Fourier coefficients, corresponding to the amplitudes of the histogram bars.

That is, if the first bar's amplitude is applied to a tone of 440 Hz (the musical note "A"), the second bar will be the amplitude of its first harmonic, 880 Hz, and so on. The superposition of these pure tones will have the same fundamental frequency but will exhibit a different timbre depending on the amplitudes of the bars.

The sonification of the written text must be comparatively analyzed for the two entropy coding techniques: fixed- and variable-length coding. Fixed-length methods will not introduce any additional compression on the given sequence based on redundancy (this is the case of the resulting bars in Figure 1(a) and (b). On the contrary, variable-length coding (exemplified by Huffman [30]) will benefit from this redundancy, reducing the coded sequence length by assigning shorter-length codes to the most repeated (in someway redundant) letters and longer codes to the less repeated ones. The amplitudes of the histogram bars in Figure 1(c) (variable length) will be flatter than those in Figure 1(a) as well as in Figure 1(d) versus Figure 1(b), after reordering. By reproducing these two cases, in accordance with the aforementioned sonification rule, it should be possible to distinguish the sounds of the two coders.

In the second part of this exercise, students must reduce redundancy by randomly erasing a quarter of the most repeated letters, "A" and "E" (the synthesis phase). This random erasure of symbols is similar to the compressing effect that perceptual coding has on data, which students are addressing in this topic. After erasing these extra, redundant letters, it is expected that the meaning of the text will still be understandable, which is something to be assessed by each student. Based on the histogram of the reduced (compressed) text, students must produce the new related sound following the same steps as before and devise comparative concluding remarks.

The reduction of the redundancy produced by an effective coder (in comparison to a less-effective one) could also be perceived as a different instrument.

The students will also develop their DSP programming skills through this exercise. They must program the coding of text using fixed- and variable-length coding techniques and then produce musical notes according to the histograms. Implicitly, students will be operating with the digital representation of the reproduced musical notes, which, in turn, corresponds to an understanding of the digital sampling principles used to produce sound.

Exercise 3: A meeting between Pollock and Picasso through the concept of redundancy

This exercise attempts to illustrate the concept of information entropy in the perception of images. At the same time, we aim to develop technical skills related to the use and manipulation of the bidimensional discrete cosine transform (DCT). The perception of entropy may be naturally inferred from images. To illustrate this, we have designed this exercise based on paintings by Picasso and Pollock [Figure 2(a) and (b), respectively]. The upper-left corner of Pollock's painting has been selected to obtain a picture that is the same size as Picasso's. Broadly speaking, the frame by Pollock seems to exhibit a higher entropy than that by Picasso, with regard to the variety of colors, intensities, and patterns used.

Again this exercise is divided into two parts. The first part drives the students to comparatively analyze the entropy present in Figure 2(a) and (b), based on which the analytical results are contrastingly considered regarding the perception of both images (analysis phase). In the second part, students

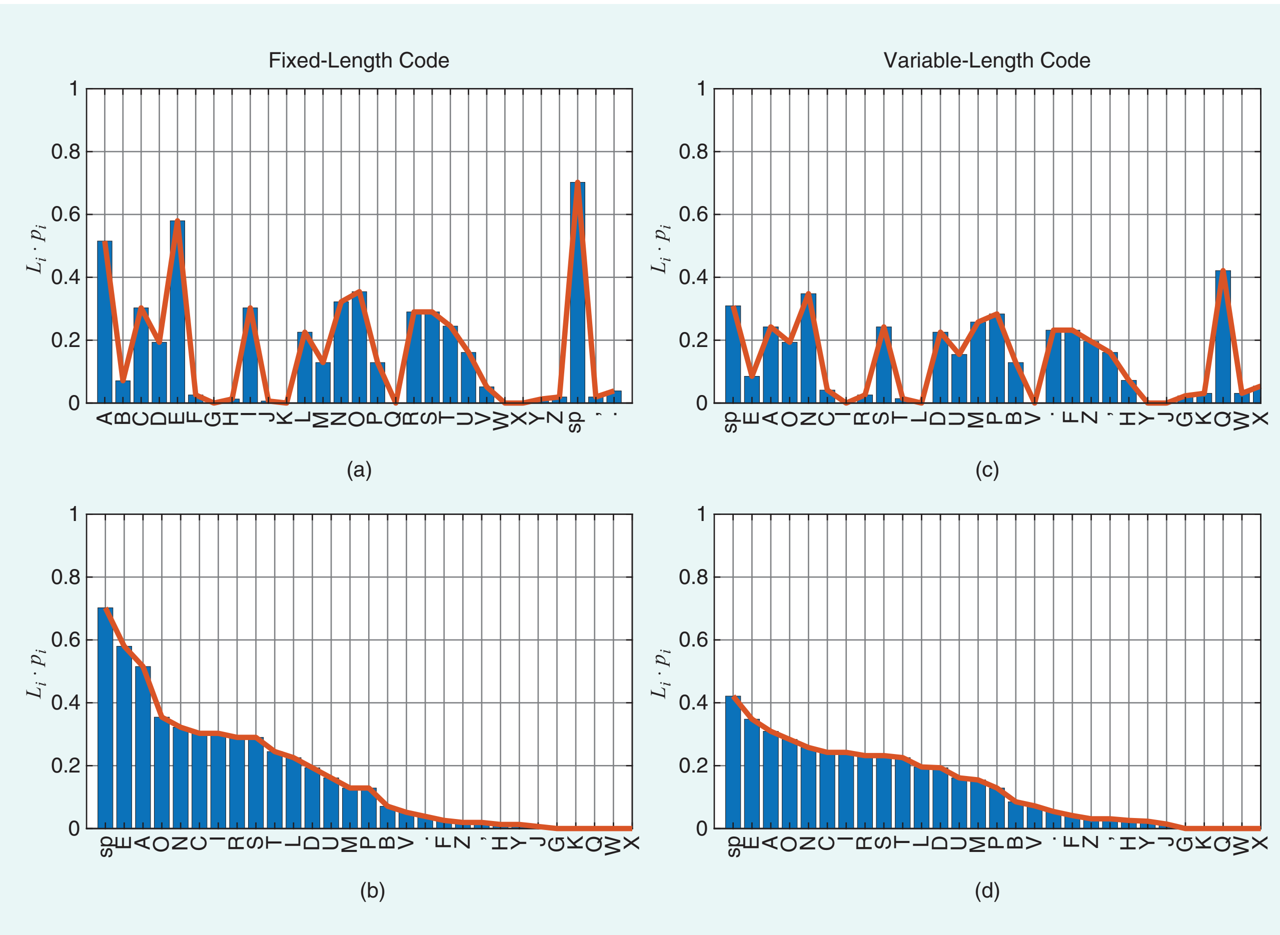

FIGURE 1. A histogram plot of $L_i \cdot p_i$ for each of the letters of a given text. (a) The alphabetical order (fixed length), (b) the descending order (fixed length), (c) the alphabetical order (Huffman), and (d) the descending order (Huffman).

are required to produce a new image (synthesis phase). This is derived by a combination of both paintings in the frequency domain through the use of the DCT. There, the entropy present in the frame by Picasso is increased to meet Pollock's entropy. The resulting frame is illustrated in Figure 2(c), where Pollock's style of painting, included in the painting by Picasso, produces an increased perceived entropy.

Specifically, in the first part of the exercise, we suggest the representation of the image as well as the computation of the entropy of the image through the use of the DCT (analysis phase). When generating the proper codes on MATLAB, the student must address the following items:

1) Obtain the matrix representation of these images and depict them in the proper plots.
2) Apply the DCT transform to each image in blocks of 8 × 8 pixels.
3) Compute the histogram of the frequency components of the image based on the obtained DCT for each block, and plot these as a 1D graph.
4) Infer conclusions regarding the entropy of the frequency components.

By solving this exercise, students are encouraged to develop a feeling for the redundancy and entropy in the frequency domain while interacting with the image and analyzing the differences between the two frames. It is worth noting that the translation of such intuitions from the spatial domain to the frequency domain is conceptually quite complex, and the fact that the students have to reach it using their own reasoning reinforces the persistence of this extremely important notion in their memories.

The second part of this exercise is related to the production of images based on the perception of entropy (synthesis phase). In this direction, Pollock's chaotic style will be incorporated into Picasso's painting in the frequency domain through the DCT. Both styles will be mixed to devise a new frame, where the entropy of Picasso's *The Blue Cup* will be increased through the replacement of its higher-frequency portion by content obtained from the painting by Pollock. The resulting operation will produce frames similar to the ones depicted in Figure 2(c), where both styles are mixed, and thereby comparatively illustrate the impact of increased entropy.

To obtain the resulting frame in Figure 2(c), students must abide by the following steps:

1) Obtain the DCT representation for both frames in Figure 2(a) and (b) in a block resolution of 8 × 8 pixels.
2) Replace the higher-frequency portion of the obtained DCT's blocks from the painting by Picasso with the DCT content of the painting by Pollock.
3) Apply the inverse DCT to the resulting operation.

The second step is applied as explained in Figure 3. The DCT blocks for each frame are obtained as indicated in Figure 3(a) and (c). Following this, a mask is defined, as depicted in Figure 3(b), to replace the higher-frequency portion of the DCT blocks from the Picasso painting with the DCT from Pollock's painting through the operation $\mathrm{DCT}_{\mathrm{Pic\&Poll}} = \mathrm{DCT}_{\mathrm{PIC}} \circ \mathrm{M} + \mathrm{DCT}_{\mathrm{Poll}} \circ (1 - \mathrm{M})$, where "∘" represents the Hadamard, or element-wise, product and $\mathrm{DCT}_{\mathrm{Pic\&Poll}}$, $\mathrm{DCT}_{\mathrm{Pic}}$, $\mathrm{DCT}_{\mathrm{Poll}}$, and M are 8 × 8-square matrices representing the DCT transforms of the resulting operation, Picasso's frame, Pollock's frame, and a binary mask, respectively. After applying this linear operation, the resulting DCT is obtained, as depicted in Figure 3(d).

As a result, when including the higher-frequency content from the painting by Pollock in the painting by Picasso, the resulting entropy will be increased. As depicted in Figure 3(c), the synthesized frame will exhibit much more randomness than will the original frame in Figure 3(a), but it will still show the composition by Picasso. In some respects, we are meeting both painters through the use of the DCT to ignite a discussion about entropy and spatial and frequency duality thorough a comparison of the results. Finally, students will develop their DSP technical skills concerning image processing when implementing and analyzing the results of the direct and inverse DCT transforms.

(a) (b) (c)

FIGURE 2. An illustrative example of a painter's and of mixed styles. (a) Pablo Picasso's *The Blue Cup* (1902). (b) Jackson Pollock's *Number 31* (1950). (c) The mixed style of Pollock and Picasso.

Quantitative and qualitative evaluation

Two types of evaluations have been carried out: first, a quantitative questionnaire about the concepts that have been acquired, and second, a qualitative questionnaire used to gather the students' opinions.

The quantitative questionnaire is the final exam in which a combination of theoretical and practical questions aim to assess the student's performance in both types of skill. The same type of exam was carried out in previous editions of this course. Due to the COVID-19 pandemic, the examination in 2020 had to be conducted remotely, and some adaptations had to be made to the content. Specifically, it was a synchronous test carried out using the assessment functions of the institutional Moodle-based learning platform. Certain measures designed to discourage answer sharing (e.g., the randomization of questions, sequentiality, and adequate time limitation) were implemented. In addition, the exam was adapted to allow students to consult their books, notes, or other sources. Surveillance was implemented using Google Meet. The students were allowed to retake the exam only once, approximately one month later.

We compared the results of the present questionnaires used with the average of the corresponding questionnaire in the last five editions of the course. This is the maximum number of editions we are allowed to retain data from according to the General Data Protection Regulation enforced in Europe. Our results show quite an improvement on the scores: a median of 56.71 and a mean of 56.64 with 23 students taking the exam versus a median of 48.86 and a mean of 47.54 with 20.2 students taking the exam averaged over the previous five years. All of the scores were higher than 100. The maximum score is 100. The standard deviations of the five-year average were 6.89, 5.22, and 2.58 for the median, mean, and number of students, respectively. Although the results are highly positive, we need to be cautious about the relatively low number of students and the fact that COVID-19 has certainly had impacts that are very difficult to quantify.

In some respects, we are meeting both painters through the use of the DCT to ignite a discussion about entropy and spatial and frequency duality thorough a comparison of the results.

The results of the retake exam were not as promising; however, the number of students taking it was sensibly lower and their abilities are different (since the students are not able to retake the exam if they have already passed it) a median of 47.9 and a mean of 51.3 with seven students taking the examination versus a median of 51.66 and a mean of 51.6 with 8.8 students averaged over the previous five years. The standard deviations of the five-year average were 7.6, 9.8, and 3.56 for the median, mean, and number of students, respectively.

The qualitative test contained the following six statements, with several options for each of the three exercises (an average of 13.33 students answered this survey):

1) The activity was (very interesting, interesting, barely interesting, not interesting); 91.07% of the students chose one of the first two options.
2) The activity was (very complex, complex, barely complex, not complex); 92.16% chose one of the first two options.
3) My previous knowledge on the contents was (enough, not enough) to complete the lab exercise; half of the students answered "enough."
4) The information provided for the lab exercise was (very adequate, adequate, scarce, very scarce); 64% chose one of the first two options.
5) The instructor's explanations (greatly helped me, barely helped me); 78.47% of the students chose "greatly helped me," and 21.53% selected "barely helped me."
6) The lab exercise helped me (very much, somewhat, very little, not at all) to understand key concepts; 80.01% chose one of the first two options. The complete breakdown of the answers can be found in Table 1.

From this we can conclude that the majority of the students found the exercises interesting and felt that they helped them to understand the key concepts targeted. However, they also found the exercises somewhat complex (because of the efforts required to derive meaningful results), and some of them would have liked further guidance material. This underspecification was a feature that we sought so as to keep the instructions open enough to foster creativity; this may

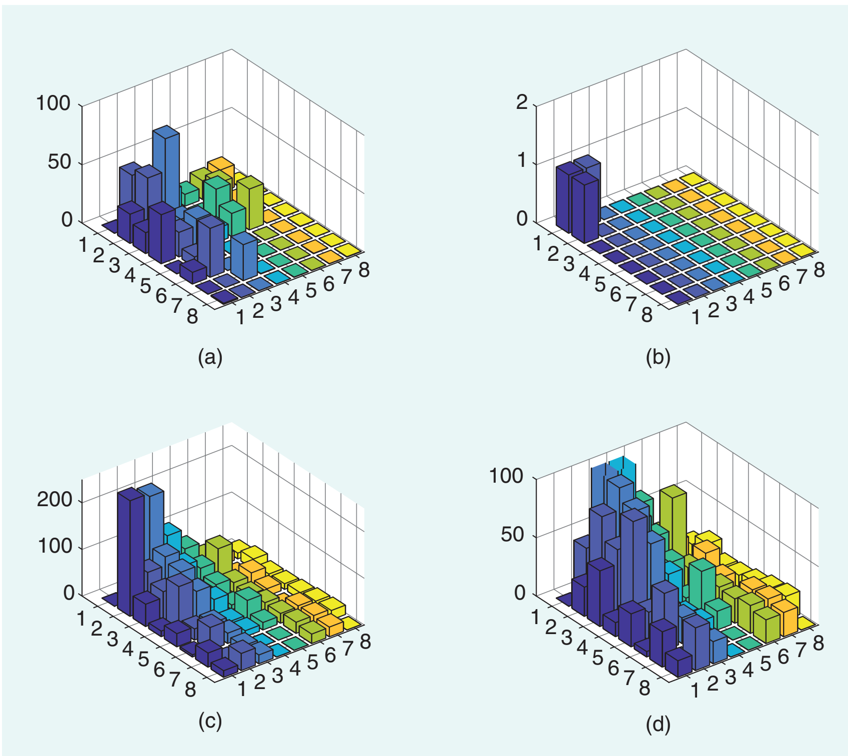

FIGURE 3. An illustrative example of a painter's and of mixed styles. (a) Pablo Picasso's *The Blue Cup* (using a discrete cosine transform) (1902). (b) A sample mask (using a discrete cosine transform). (c) Jackson Pollock's Number 31 (1950). (d) The mixed style of Pollock and Picasso (using a discrete cosine transform).

Table 1. The qualitative tests for exercises 1–3 (the averaged percentage and standard deviation). The particular transcriptions of the options for each of the questions are described in the text. It is worth noting that questions 3 and 5 are binary.

Q	Option 1	Option 2	Option 3	Option 4
1	(21.29, 18.74)	(69.78, 15.8)	(4.76, 8.25)	(4.17, 7.22)
2	(14.46, 20.17)	(77.7, 19.73)	(7.84, 13.58)	(0, 0)
3	(49.02, 1.7)	(50.98, 1.7)	—	—
4	(14.39, 8.36)	(49.61, 13.57)	(31.65, 15.92)	(2.38, 4.12)
5	(78.47, 8.22)	(13.03, 11.44)	—	—
6	(33.37, 10.5)	(46.68, 3.59)	(11.48,15.09)	(6.55, 6.27)

Table 2. The qualitative tests for overall satisfaction (percentage) answered by 14 students. The particular transcriptions of the options for each of the questions are described in the text. It is worth noting that question 4 is binary.

Q	Option 1	Option 2	Option 3	Option 4
1	28.57	42.86	28.57	0
2	35.71	50	14.29	0
3	28.57	50	14.29	7.14
4	92.56	7.14	—	—

be the reason for why the instructor's indications were so highly appreciated.

An overall satisfaction questionnaire was also appended to the third qualitative test and contained the following four statements:

1) The activities were (very, somewhat, barely, not at all) original and innovative; 71.43% of the students chose one of the first two options.
2) The activities were (very much, somewhat, barely, not at all) related to the engineering profession; 85.71% of the students chose one of the first two options.
3) The activities prepared me for critical thinking and reflection (very much, somewhat, barely, not at all); 78.57% of the students chose one of the first two options.
4) The activities challenged me (yes, no); 92.86% of the students chose "yes." The complete breakdown of the answers can be found in Table 2.

We would like to highlight here that nearly all of the students found the activities challenging, which was one of our main objectives. They also very clearly perceived the relationship the exercises had with professional engineering. The scores on originality and critical thinking were also very positive but slightly lower.

Discussion

We used the proposed exercises detailed in the previous section to raise our student's interest in the liberal arts. The students employed narrative techniques to produce their own scripts in exercise 1, they made adjustments so that the sound produced had better sonority or acoustic properties in exercise 2, and they appreciated the composition of paintings in exercise 3.

The exercises were designed to combine both creative and more-guided learning processes. For example, exercise 1 required that students participate actively when writing and recording their own ideas; we expected them to express their concerns and proposals about the topics under analysis. Following this, measuring the information entropy metric on a given text is a more standard procedure, during which they should follow straightforward specifications. The students should also become critical about the concept itself and its usefulness, which was verified empirically a posteriori once they had actively removed irrelevant content from their own excerpts of text. They proved themselves able to check that their text remains understandable, but observe at the same time, its length and entropy change, illustrating the perceptual coding strategies.

Exercise 2 illustrates the concept of redundancy in terms of sound perception. The students may appreciate the concept of redundancy when they try to understand their own text after randomly erasing the most repeated letters. In this regard, redundancy is portrayed as the part of an object that may be removed without losing its descriptive ability. This idea is translated into sounds via sonification schemes, encouraging a perception of redundancy by listening.

Finally, exercise 3 produces a visual representation of entropy to be contrasted with its value (analytically derived). Here we can perceive comparatively higher and lower values of entropy in the graphical content of paintings. The students also produce a mixed style of painters by incorporating additional entropy from one frame to the other. When trying to balance both frames from an information entropy standpoint, the higher-frequency portion of one frame is replaced by the other. This mixture of analytical operations and the perception of the resulting images naturally facilitates the assimilation of the concepts of entropy and redundancy.

At the same time, students discuss the core concepts of engineering and develop technical skills related to DSP subjects. In exercise 1, students reflect their understanding of the concept of entropy without the support of any mathematical formula, which gives a broaden idea of their comprehension of the topic. Exercise 2 reflects the concept of redundancy in terms of sound perception, while exercise 3 does the same through image. Considering the technical skills, hands-on activities involving core DSP techniques (sampling, representation, and transformation) are conducted when undertaken these exercises. The students develop the proper codes to read, store, analyze, process, and produce multimedia content (text, sound, and image). In addition, students are introduced to the basic implementation of the most common standards used to transmit multimedia content. The application of encoders and signal processing through the DCT are the basic elements of JPEG, MPEG, and H.26x standards.

Essentially, the proposed exercises are deliberately underdefined to give the students the chance to provide (and discuss) their own definitions of how redundancy should be measured for text, music, or images. In this respect, a connection with nonengineering departments may support the introduction of some basic liberal arts concepts. For instance, the teacher may introduce the elementary techniques used to develop narrative skills, wherein students are introduced to scriptwriting to support the solutions to exercises 1 and 2. A discussion of how to express the same concept but with less text and still preserving the meaning can also be carried out when developing solutions to exercise 1. Alternatively, the teacher may prompt a discussion about the perception of redundancy in different musical rhythms; for example, they may pose questions such as "What is more redundant, classical music by Bach or the flamenco guitar by Paco de Lucía?"; "What sound attribute will be more useful to model and reduce redundancy?"; and "Is there some correlation between redundancy and something as subjective as the quality of music?" Note that these issues will inspire a discussion even among specialists on the subject because the connections we are trying to evoke here are by no means straightforward.

The evaluation of this study has proceeded in two directions: first, a quantitative questionnaire about the concepts to be acquired, and second, a qualitative questionnaire aimed at gathering student feedback. Our results show a positive improvement in the scores, although the sample of students is rather small. The qualitative test comprised six questions related to the students' opinions of the three new exercises; the questionnaire covered students' opinions regarding thought inspiration, the complexity of exercises, supporting theoretical background, and the completion of learning objectives. Four final questions concerning their overall opinions were appended to the last questionnaire, and the response showed a very positive attitude toward the new materials.

Future directions: Further DSP connections with the arts

Our methodology relies on the identification of core concepts in the targeted course and their connection with the arts as the main driving force. Here we suggest connections with other DSP core concepts. For example, the fundamentals of the digitalization of analog signals is mediated by Nyquist's sampling theorem and its requirements for valid representations of continuous-time signals and systems in a discrete-time domain. As is well known, the sampling theorem allows for the acquisition of equivalent discrete sequences from their analog counterparts. Similarly, discrete time-domain operations, such as discrete infinite-impulse-response filters and, most significantly, their discrete Fourier transforms, have output–input relationships equivalent to analog systems [1]. Both sampling and filtering, which allow the reproduction of analog signals and systems by their discrete-time counterparts, together with linear transformations such as Fourier's (which bring a very convenient duality in time and frequency processing), are general and universal concepts in discrete-time signal processing.

Several connections with the arts can be identified to illustrate, for example, the meaning and implications of the sampling theorem. The impact of the sampling period on sound reproduction or the representation of images can be made (perceptually) very noticeable in several creative ways. Besides the obvious (but also interesting) dependencies that can be observed by down- or upsampling the signals, musical elements such as rhythm and style, instrument tessituras, or vocal ranges can be exploited in relation to the sampling periods needed to process them in the discrete-time domain. For instance, Franz Schubert's *Ave Maria* (Op. 52 No. 6), interpreted on a violin will provide a sharp contrast when compared to the main theme from the *Interstellar* soundtrack by Hans Zimmer, covered by a violoncello because of the higher-frequency tones used in the former. Similarly, the styles from certain art movements (for example, pointillism techniques) will require more reduced sampling periods than will the blue period paintings by Picasso. Note that these last notions and their connections to the sampling frequency are more nuanced and conducive to discussion than purely engineering-based considerations.

In a similar fashion, the notion of filtering as applied to sound (a 1D filter) or to a given frame or photograph (a 2D filter) will modify our perception of the musical piece or the paint depending on the filter's specifications. The impact of the design of the filters (i.e., low or high pass, and so on) can be discussed in connection with the instrumentation of a musical piece. Through the use of 2D filters can compare the results of smoothing and sharpening effects on visual pieces of art from two different art schools or periods. Again, for instance, smoothing effects will be more pronounced on pointillism than on the blue paintings by Picasso.

Conclusions

This ongoing research presented in this article introduced the liberal arts as a way to improve the teaching of abstract concepts in engineering. By making innovative use of writing, sonification, and image production techniques, we offered lab activities that interconnect the attributes of text, audio, and image with an active use of liberal arts concepts. We believe that this approach reinforces an understanding of key concepts and reduces student perceptions of their difficulty, not only through their analytical definition but also through the feelings or vibes they convey. This proposal also leads to the integration of nonengineering departments to design mixed curricula where both schools are integrated. Our evaluations indicated positive outcomes regarding the learning objectives and appreciation of the course. Future directions include devising original exercises to cover new topics and key concepts and broadening the number and type of students involved to test their impact.

Acknowledgments

This research is part of the 17th Call to Support Experiences on Innovative Teaching at Universidad Carlos III de Madrid (UC3M), Spain, in 2019–2020 and partially funded by the Spanish Government-MinECo project TEC2017-84395-P. This project gathered members from UC3M, the Havana University of Technology, José Antonio Echeverría, Cuba, and Technische Universität Berlin, Germany.

Authors

Jorge Torres Gómez (torres-gomez@ccs-labs.org) received his Ph.D. degree from the Havana University of Technology José

Antonio Echeverría (CUJAE), Havana, Cuba, in 2015. He is currently with the School of Electrical Engineering and Computer Science, Technische Universität (TU) Berlin, Berlin, 10587, Germany. From 2008 to 2018, he lectured at the School of Telecommunications and Electronics, CUJAE. He has been with the Department of Signal Theory and Communications, Universidad Carlos III de Madrid, Madrid, Spain, as a guest lecturer, and with TU Chemnitz as a postdoctoral researcher. He is with the Committee of Teaching and Professional Activities in the IEEE German Chapter Executive Committee. His research interests include digital signal processing, software-defined radio, and education. He is a Member of IEEE.

Antonio Rodríguez-Hidalgo (arodriguez@tsc.uc3m.es) received his Ph.D. degree in multimedia and communications from Universidad Carlos III de Madrid (UC3M), Madrid, 28911, Spain, in 2019, where he is currently with the Department of Signal Theory and Communications. He was a visiting researcher with the Chait-Lab, the Ear Institute, University College London, in 2017. He was also a postdoctoral researcher at UC3M, where his research focused on autoencoders applied to audio signals.

Yannelys Virginia Jerez Naranjo (yanny@tesla.cujae.edu.cu) received her master's degree in technology applied to educational processes from the Havana University of Technology José Antonio Echeverría (CUJAE), Havana, 10390, Cuba, where she is currently the director of scientific and technical information. She is also an assistant professor at CUJAE's Reference Center for Advanced Education, specializing in infotechnologies. She participates in projects related to the integration of information and communications technology in higher-education processes and teaches undergraduate and graduate students.

Carmen Peláez-Moreno (carmen@tsc.uc3m.es) received her Ph.D. degree from Universidad Carlos III de Madrid (UC3M), Madrid, 28911, Spain, where she is currently an associate professor in the Department of Signal Theory and Communications. She collaborates with the Spanish National Research Agency and has earned the Full Professorship National Habilitation. She was a visiting researcher at the University of Westminster, London, United Kingdom; the University of Strathclyde, Glasgow, United Kingdom; the University of Trento, Trento, Italy; and the International Computer Science Institute, Berkeley, California, USA. Her research interests include speech recognition and perception, affective computing, machine learning and data analysis, information theory, and education. She is a Member of IEEE.

References

[1] A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing* (Prentice-Hall Signal Processing Series). London: Pearson, 2010.

[2] C. P. Snow, *The Two Cultures and the Scientific Revolution*. Mansfield Center, CT: Martino Fine Books, 2013.

[3] J. D. Klein and R. Balmer, "Engineering, liberal arts, and technological literacy in higher education," *IEEE Technol. Soc. Mag.*, vol. 26, no. 4, pp. 23–28, 2007. doi: 10.1109/MTS.2007.911066.

[4] L. L. Bucciarelli and D. E. Drew, "Liberal studies in engineering: A design plan," *Eng. Stud.*, vol. 7, no. 2-3, pp. 103–122, Sept. 2015. doi: 10.1080/19378629.2015.1077253.

[5] M. D. Shuster, "The arts and engineering [Focus on Education]," *IEEE Control Syst. Mag.*, vol. 28, no. 4, pp. 96–98, Aug. 2008. doi: 10.1109/MCS.2008.924881.

[6] J. D. Klein, C. Traver, S. Raucci, and S. A. Jones, "Special session—inventive courses: Blending engineering and the liberal arts," in *Proc. 39th IEEE Front. Educ. Conf.*, Oct. 2009, pp. 1–2. doi: 10.1109/FIE.2009.5350505.

[7] D. A. Sousa and T. Pilecki, *From STEM to STEAM: Using Brain-Compatible Strategies to Integrate the Arts*. Thousand Oaks, CA: Corwin Press, 2013.

[8] L. Colucci-Gray, P. Burnard, D. Gray, and C. Cooke, "A critical review of STEAM (science, technology, engineering, arts, and mathematics)," in Oxford Research Encyclopedia of Education, P. Thomson, Ed. London, U.K.: Oxford University Press, 2019, pp. 1–26.

[9] E. Watson, "STEM or STEAM?: The critical role of arts in technology education (and the critical role of art in technology)," *Irish J. Acad. Pract.*, vol. 8, no. 1, p. 8, 2020.doi: 10.21427/eqzb-vb42.

[10] L. Catterall, "A brief history of STEM and STEAM from an inadvertent insider," *STEAM J.*, vol. 3, no. 1, p. 5, 2017. doi: 10.5642/steam.20170301.05.

[11] C. Liao, "From interdisciplinary to transdisciplinary: An arts-integrated approach to STEAM education," *Art Educ.*, vol. 69, no. 6, pp. 44–49, 2016. doi: 10.1080/00043125.2016.1224873.

[12] A. Oner, S. Nite, R. Capraro, and M. Capraro, "From STEM to STEAM: Students' beliefs about the use of their creativity," *STEAM*, vol. 2, no. 2, p. 6, 2016. doi: 10.5642/steam.20160202.06.

[13] "Science and engineering indicators," National Science Foundation, Alexandria, VA, 2018. [Online]. Available: https://nsf.gov/statistics/2018/nsb20181/

[14] J. R. Cimpian, T. H. Kim, and Z. T. McDermott, "Understanding persistent gender gaps in STEM," *Science*, vol. 368, no. 6497, pp. 1317–1319, 2020. doi: 10.1126/science.aba7377.

[15] K. Panetta, "The art of engineering [letter from the editor]," *IEEE Women Eng. Mag.*, vol. 13, no. 2, pp. 2–3, 2019. doi: 10.1109/MWIE.2019.2938929.

[16] N. Salim, "The 'IT' girls: Inspiring girls in science, technology, engineering, art, and math [pipelining: attractive programs for women]," *IEEE Women Eng. Mag.*, vol. 6, no. 2, pp. 38–40, 2012. doi: 10.1109/MWIE.2012.2215263.

[17] "Pasaporte Cultural," Universidad Carlos III de Madrid. Accessed Feb. 24, 2019. [Online]. Available: https://www.uc3m.es/ss/Satellite/Cultura/es/TextoMixta/1371230632607/

[18] L. Bucciarelli and D. E. Drew, "Breaking Boundaries with Liberal Studies in Engineering," in *The Future of Engineering Education is Interdisciplinary*. P. Kapranos, Ed. Evanston, IL: Routledge, Apr. 2018, pp. 133–146.

[19] H. M. Jawad, S. Tout, M. Abualkibash, and Y. Xie, "Integrating art and animation in teaching computer programming for high school students experimental study," in *Proc. IEEE Int. Conf. Electro/Inform. Technol. (EIT)*, May 2018, pp. 0311–0317. doi: 10.1109/EIT.2018.8500178.

[20] A. Barmpoutis, "Integrating algebra, geometry, music, 3D art, and technology using emoticoding," in *Proc. IEEE Integrated STEM Educ. Conf. (ISEC)*, Mar. 2018, pp. 30–33. doi: 10.1109/ISECon.2018.8340500.

[21] Q. Fengmei and X. Zhang, "Construct a curriculum system of high vocational education combined technology with arts to implement the cultivation of innovative talents of webdesigner," in *Proc. 2nd Int. Symp. Instrumentation Measurement, Sensor Netw. Autom. (IMSNA)*, Dec. 2013, pp. 83–88. doi: 10.1109/IMSNA.2013.6742823.

[22] N. A. Abu, M. D. Hassan, and S. Sahib, "Mathematical animations: The art of teaching," in *Proc. 31st Annu. Frontiers Educ. Conf. Impact Eng. Sci. Educ. Conf.*, Oct. 2001, vol. 3, pp. S1C–S10. doi: 10.1109/FIE.2001.963987.

[23] P. Ekman, "An argument for basic emotions," *Cognit. Emotion*, vol. 6, no. 3–4, pp. 169–200, 1992. doi: 10.1080/02699939208411068.

[24] G. Brod, M. Hasselhorn, and S. A. Bunge, "When generating a prediction boosts learning: The element of surprise," *Learn. Instruction*, vol. 55, pp. 22–31, June 2018. doi: 10.1016/j.learninstruc.2018.01.013.

[25] R. N. Henson and P. Gagnepain, "Predictive, interactive multiple memory systems," *Hippocampus*, vol. 20, no. 11, pp. 1315–1326, 2010. doi: 10.1002/hipo.20857.

[26] A. Greve, E. Cooper, A. Kaula, M. Anderson, and R. Henson, "Does prediction error drive one-shot declarative learning?" *J. Memory Lang.*, vol. 94, pp. 149–165, June 2017. doi: 10.1016/j.jml.2016.11.001.

[27] R. E. Mayer, "Multimedia learning," *Psychol. Learn. Motivation*, vol. 41, pp. 85–139, Jan. 2002. doi: 10.1016/S0079-7421(02)80005-6.

[28] J. D. Gibson, T. Berger, T. Lookabaugh, R. Baker, and D. Lindbergh, *Digital Compression for Multimedia: Principles and Standards*. San Mateo, CA: Morgan Kaufmann, 2011.

[29] S. Misra, M. Reisslein, and G. Xue, "A survey of multimedia streaming in wireless sensor networks," *IEEE Commun. Surveys Tuts.*, vol. 10, no. 4, pp. 18–39, 2008. doi: 10.1109/SURV.2008.080404.

[30] D. A. Huffman, "A method for the construction of minimum-redundancy codes," *Proc. IRE*, vol. 40, no. 9, pp. 1098–1101, 1952. doi: 10.1109/JRPROC.1952.273898.

SP

Michael Gerstmair, Martin Gschwandtner, Rainer Findenig,
Oliver Lang, Alexander Melzer, and Mario Huemer

Miniaturized Advanced Driver Assistance Systems

A low-cost educational platform for advanced driver assistance systems and autonomous driving

©SHUTTERSTOCK.COM/LUCKYSTEP

One exciting challenge for future scientists and engineers will be to contribute to the development and introduction of self-driving cars. Under the objective of making young talent attending engineering-focused secondary schools and universities enthusiastic about this highly innovative topic, we have developed an educational platform for advanced driver assistance systems (ADASs) and autonomous driving. Our platform is designed for a miniaturized environment with low cost and scalable complexity, simultaneously providing all of the sensors required to automate a vehicle.

To make Miniaturized ADAS suitable for a wide range of students, we integrated our platform into a LEGO train as well as radio-controlled (RC) model cars and developed an assignment catalog that includes several levels of difficulty. The main focus of this catalog is to convey the basic functionality of the sensors as well as the underlying signal processing techniques in a hands-on manner. Depending on the educational level and desired learning objectives, interested lecturers have the opportunity to select arbitrary modules to build up practical courses in accordance with their preferences.

In this article, we provide insight into this multidisciplinary educational project developed together with students for students. In addition, we introduce the most important sensors currently used in ADASs and self-driving vehicles. To illustrate the broad spectrum and multitude of possibilities of this project, we further give a rough overview of the applied signal processing strategies and highlight the weaknesses and strengths of the individual sensors. To give an impression of what practical courses for university-level students based on Miniaturized ADAS might look like, we finally present an exemplary assignment catalog for a radar-based ADAS. The reason for focusing on radar is that, to the best of our knowledge, Miniaturized ADAS is the only educational platform supporting frequency-modulated continuous wave (FMCW) radar sensors.

Introduction

Nothing is constant but change—from a technical and economic point of view, this statement has been especially true

Digital Object Identifier 10.1109/MSP.2021.3051939
Date of current version: 28 April 2021

since the first industrial revolution at the end of the 18th century. This point in time marked the transition from handcrafted production to machine manufacturing, which led to the mass production of products for the first time in human history. While the first machines, like the loom, had to be operated manually, subsequent engines were driven by water and steam power. This progress led to the first effectively working steam locomotive—or rather, steam car—designed and built by William Murdock and based on patents by James Watt in 1784 [1]. By establishing the electrical current at the end of the 19th century, even more automation became possible, which ultimately led to the second industrial revolution. The most common innovation from that time is probably represented by the assembly line work successfully realized by Henry Ford in 1913 [2]. This step made the automobile affordable for new customers and therefore suitable for the mass market.

The third industrial revolution, or digital revolution, typically dated to the middle of the 20th century, embodies the conversion from mechanical and analog electronics to digital electronics. This technological improvement enabled the integration of ADASs, such as adaptive cruise control (ACC) or the lane keep assistant (LKA) into high-class cars around the turn of the millennium [3]. The development of these safety features significantly decreased the number of injuries and fatalities on the road [4]. Further, it seems obvious that these numbers can be reduced even more by introducing self-driving cars, which are an essential part of the fourth industrial revolution, also denoted as Industry 4.0.

These technical advancements are accompanied by the emergence of a huge number of new research areas and possibilities in the field of signal processing and machine learning. Obviously, it is essential to excite young researchers and engineers to work in these areas. But beyond that, it should also be of the utmost importance to adequately educate young students to be well prepared for the challenges ahead. Due to the fact that technological progress and the associated changes have become significantly faster, we should adapt our educational system accordingly.

Besides imparting important well-founded fundamental knowledge, it is also necessary to promote students' creativity and problem-solving competencies on the system level. According to David Kolb's experiential learning model, this may be efficiently done in a hands-on manner in which the students receive direct feedback and can reflect on their developed and implemented solutions [5]. Motivated by these considerations, we started a collaboration between Infineon and the Institute of Signal Processing at Johannes Kepler University (JKU) Linz, with the aim to develop an ADAS in a miniaturized environment. Extending this cooperation to various educational institutions with different focuses finally led to this project.

In the following, we give insight into this multidisciplinary educational project, with a special focus on signal processing, and describe how it can be integrated into engineering-focused secondary school courses and university lectures or courses.

The Miniaturized ADAS project

The idea to develop self-driving model vehicles has already been realized by different institutions. Typically, such a development goes along with competitions such as the Carolo-Cup [6], the NXP-Cup in Europe, the Middle East, and Africa (EMEA) [7], or the Audi Autonomous Driving Cup (AADC) [8]. The objective of these contests is to develop a self-driving vehicle that is able to master various tasks within the time frame of one year. Although the mentioned competitions seem to be similar at first glance, they substantially differ from each other. The Carolo-Cup, organized by Technische Universität Braunschweig, requests that contributing students develop, build, and demonstrate a cost- and energy-efficient 1:10 model of an automated vehicle from scratch [9]. The problem definition of the AADC, on the other hand, is fundamentally different. Here, a 1:8 scale replica of an Audi Q2 is provided to student teams in the course of the competition. This means that every team is forced to use the same hardware to master a given obstacle course [8]. As a third variant, the current edition of the NXP-Cup EMEA allows the students to either use an existing NXP Cup car kit, acquire a commercial kit, or even build their own kit [7].

By comparing the mentioned competitions, we came to the following conclusions. Providing an already prefabricated hardware platform—including a tested industry-standard software environment like that for the AADC—enables the students to focus on the design and implementation of the decision logic. This saves time since there is no necessity to deal with the integration of the sensors. On the other hand, a disadvantage of this powerful platform is its very high price. In contrast, allowing the students a high degree of freedom by letting them build the vehicle from scratch significantly widens the required tasks and expertise. This, in further consequence, will lead to larger teams and also demand management skills in addition to technical knowledge. However, from our experience, it is very challenging to build a smart and robust high-quality vehicle within only one year.

Therefore, based on two primary objectives, the idea of Miniaturized ADAS is to develop an affordable and highly configurable platform together with students for students. First, this platform is used to automate different kinds of model vehicles providing in-depth guidance to a small group of students. This is typically done in the course of internships and is often combined with appropriate examination projects or theses. Second, we used the developed demonstrators to prepare workshops for colleges and practical courses or lectures for universities. Furthermore, interested students attending these courses get the opportunity to become a part of Miniaturized ADAS and contribute to improving our platform and educational applications. To this end, we collaborate with various types of educational institutions.

The entry point into Miniaturized ADAS is represented by engineering-focused secondary schools. These five-year courses start with grade nine and are typically denoted as a higher technical school, or Höhere Technische Lehranstalten (HTLs), in Austria [10]. The assignments for these HTL workshops, presented in [11], are matched to the educational level of the students and aim to motivate them for science, technology, engineering, and mathematics studies. To further use our platform for teaching at the university level and reach exactly these fields of study, Miniaturized ADAS can be adapted in multiple ways.

First, to scale the demands regarding the complexity of the implemented solutions, we integrated our platform into different vehicles like cars and trains, each presenting individual requirements. This is driven by the fact that, for example, the automation of a train does not require any steering effort, whereas keeping a car in a certain lane is far more difficult. In the "Existing Hardware Platforms" section, we present a selection of the automated vehicles that are intended to be used for teaching and provide some background information about their development and related collaborations.

Second, to provide a wide range of university courses with a scalable level of difficulty, we prepared a catalog containing a variety of different assignments. This enables interested lecturers to arbitrarily assemble their courses and put the focus on specific topics of their choice. An excerpt of a potential assignment catalog for practical university courses on radar signal processing and radar-based ADAS is explained in the "Course Assignments and Learning Objectives" section. Obviously, the presented catalog can easily be extended for every available sensor of the platform. Considering how many available sensors there are as well as the issues that have to be solved to fully automate a vehicle, the number of potential courses and lectures is huge. Furthermore, by adjusting or replacing single assignments, the courses can be adapted from year to year.

In this article, we chose to present an exemplary catalog based around radar because, to the best of our knowledge, there is no platform available that offers the possibility to easily include radar sensors. Typically, most miniaturized vehicles for autonomous driving base their decisions solely on camera data [12]–[14]. More advanced vehicles like the already mentioned Audi Q2 further provide ultrasound sensors or even lidar to sense the environment [8], [15], [16]. In our opinion, radar is one of the most essential sensors, especially for educational purposes since a lot of state-of-the-art ADASs, like ACC, blind spot assistant (BSA), or autonomous emergency braking (AEB), are based on radar data. In addition, ultrasound sensors will probably be replaced by radar sensors in the future, while lidar is currently hardly ever used for ADASs in modern cars.

Available sensors and actuators

The main processing unit of Miniaturized ADAS is a Raspberry Pi minicomputer. Although there are more powerful computers with a similar form factor available, the Raspberry Pi is supported by a huge community. This lowers the entry barrier, reduces the supervision expenditures, and enables students to share their knowledge and problems with a high number of engineers and so-called makers via the Internet.

Our platform further provides a variety of sensors that can be used to sense the environment of an automated vehicle. According to [17], the most important sensors to enable autonomous driving are cameras, radar sensors, lidar, and ultrasound sensors. In addition, our platform supports a current sensor, a reed relay, a gyroscope/accelerometer, and a number of different actuators that can be used for various applications. Table 1 provides a list of all of the sensors and actuators supported by the software framework of Miniaturized ADAS. (For a detailed introduction on the software framework of Miniaturized ADAS, see [11].)

However, in the following, we specifically focus on cameras, radar, lidar, and ultrasound sensors. In the course of this overview, we provide a more detailed description of the sensor capabilities as well as their advantages and disadvantages. This, in further consequence, will exemplify the wide range of signal processing areas covered by this project. Furthermore, it will point out the necessity of advanced sensor fusion strategies for future ADASs and autonomous driving.

Cameras

Among the aforementioned sensors, camera systems are probably the most versatile systems for ADASs and autonomous driving. From an image signal processing point of view, applications range from basic lane detection algorithms up to highly sophisticated machine learning approaches for object detection and classification. The former is often done with standard signal processing techniques, like edge detectors in combination with a Hough transform. Beyond that, Kalman filters can be used to predict the further course of the road [18]. The resulting information about the roadway is typically passed on to an ADAS, like a lane departure warning system or the LKA.

However, to reliably detect and classify objects, convolutional neural network (CNN)-based single-shot detectors like You Only Look Once (YOLO) and Single-Shot Detector (SSD)

Table 1. An overview of the available sensors and actuators as well as their intended use and costs.

Sensor/Actuator	Use Case	Costs
Radar 77 GHz	Range, velocity, and angle estimation	€250
Radar 60 GHz	Range and velocity estimation	€150
Rotating scanning lidar	Range and angle estimation, path finding	€100
Time of flight 3D camera	Range and angle estimation, path finding	€300
Raspberry Pi camera	Lane/object detection	€25
Wide-angle camera	Lane/object detection	€30
Ultrasound sensor	Range estimation for short-range applications	€2
Accelerometer/gyroscope	Self-positioning	€8
Reed relay	Reference positioning	€0.5
Current/voltage sensor	Battery management	€14
H-bridge	Motor control	€3
Servo motor	Vehicle drive	€15
Intel compute stick	Machine learning tasks	€100
OLED display	Output information	€5
Fog machine	Simulate poor visibility	€60
Buzzer	Attract attention	€1
Touchscreen	Control vehicles, visualize sensor data	€70

OLED: organic LED.

are typically used. These networks are able to automatically extract features to accurately recognize and locate relevant objects [19]. The detected objects can further be used for tasks like traffic sign recognition or even camera-based AEB and ACC. The drawback of camera systems is their strong dependency on environmental conditions like weather and lighting.

To offer work packages dealing with these challenges and enabling the powerful opportunities provided by cameras, the Miniaturized ADAS project supports two different camera modules. Furthermore, we integrated the lightweight CNNs PeleeNet-SSD, YOLOv3-Tiny, and MobileNetv2-SSDLite into our software framework, which allows it to perform object detection and classification on the Raspberry Pi [20]. In Figure 1, PeleeNet was used to detect and localize the two model cars as well as a doll.

Radar

For automotive applications, typically, FMCW radar systems are used. These sensors transmit a linear frequency ramp to estimate the round-trip delay time of received object reflections to finally calculate the corresponding distances [21]. By further transmitting several consecutive ramps and utilizing a multiple-input, multiple-output antenna array, the relative velocity as well as the actual position of detected objects can be estimated. This is most commonly done by fast Fourier transform (FFT)-based evaluation techniques in combination with constant false alarm rate (CFAR) algorithms used for object detection [3], [22], [23].

Due to the fact that automotive radar sensors operate in frequency ranges from 76 to 81 GHz, they offer the unique advantage of showing good performance even in adverse lighting and bad weather conditions [24]. Therefore, radar is the preferred sensor for assistant systems like ACC, AEB, or the BSA. Automotive radar sensors are able to reach distances of up to 250 m [25]. However, due to the inherent crosstalk between the transmit (Tx) and receive (Rx) channels in the radar monolithic microwave integrated circuit (MMIC), the minimum distance of an automotive radar sensor is typically limited to about 20 cm [26].

To achieve even shorter ranges, our platform supports not only an automotive radar but also a 60-GHz radar sensor that was initially developed to perform gesture recognition for smart devices. The MMIC we used is a development version of the radar chip utilized for gesture recognition in the Google Pixel 4 [27]. This sensor is able to detect objects in the range of 1–10 m, which makes it perfectly suitable for a miniaturized environment.

Figure 2 illustrates the same scenario as that depicted in Figure 1, recorded by an automotive radar sensor. Although all of the relevant object positions as well as their velocities are visible, it can be seen that the amount of information contained in this image is limited. Furthermore, it should be noted that the used objects provide very small radar cross sections and were therefore equipped with materials with good reflection properties.

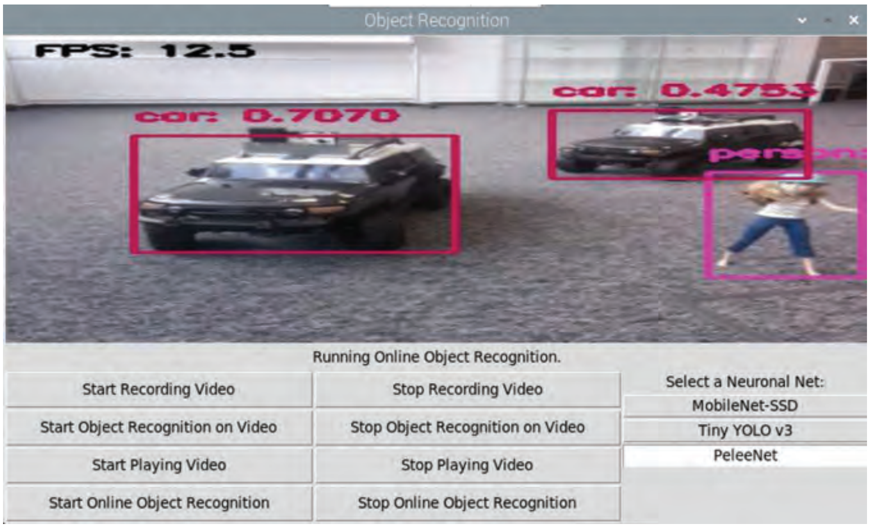

FIGURE 1. A scenario depicting two model cars as well as a doll. The objects are successfully detected by the implemented lightweight CNN executed on the Raspberry Pi.

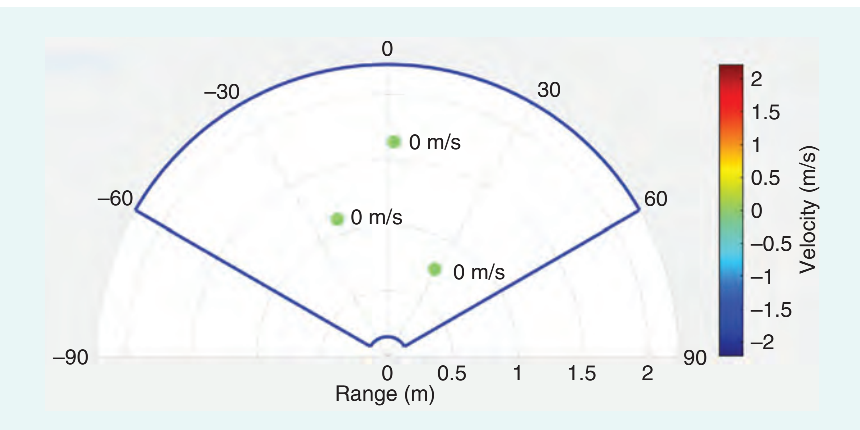

FIGURE 2. A radar image depicting two model cars and one doll equipped with very well-reflecting materials.

Lidar

In contrast to radar, a lidar system illuminates a scenery using its own light source, typically operating with wavelengths of 905–1,550 nm, to precisely analyze the reflected light and detect the angle as well as the distance of surrounding objects. In general, lidars can be divided into the categories of flash and scanning systems. A flash lidar illuminates the whole field of view simultaneously, which limits the illumination power that can be used to scan a particular spot in the environment. Scanning lidar systems, on the other hand, overcome this problem by utilizing a rotating or vibrating mirror to steer a focused laser beam in a specific direction. The most important advantage of both lidar systems is clearly the very high resolution in the angular domain. Therefore, lidar sensors are most commonly used for pathfinding in autopilot functions but could potentially also be used for systems like AEB and ACC. The drawbacks of lidar scanners are a severe degradation in performance in poor weather conditions and the limited long-term durability of the vibrating mirrors used in scanning lidar systems [28].

Miniaturized ADAS offers both of the aforementioned sensor types. The flash lidar system is represented by a so-called time-of-flight (ToF) camera [29]. These cameras are most commonly used for facial recognition in high-end smartphones like the Samsung Galaxy S10 5G or the Huawei P30 Pro. The ToF camera integrated into our platform is the CamBoard pico flex, which provides a measurement range from 0.1 to 4 m with an angular resolution of 0.27° at 45 frames/s, which is able to depict a 3D image of a scene. Beyond that, we also integrated a rotating lidar from Shenzhen EAI Technology Co. With its rotating sensor head, this lidar surveys the horizontal level in its surrounding up to a distance of 10 m. The sensor head rotates up to 12 times/s and offers measurements with an angular resolution of around 0.5°.

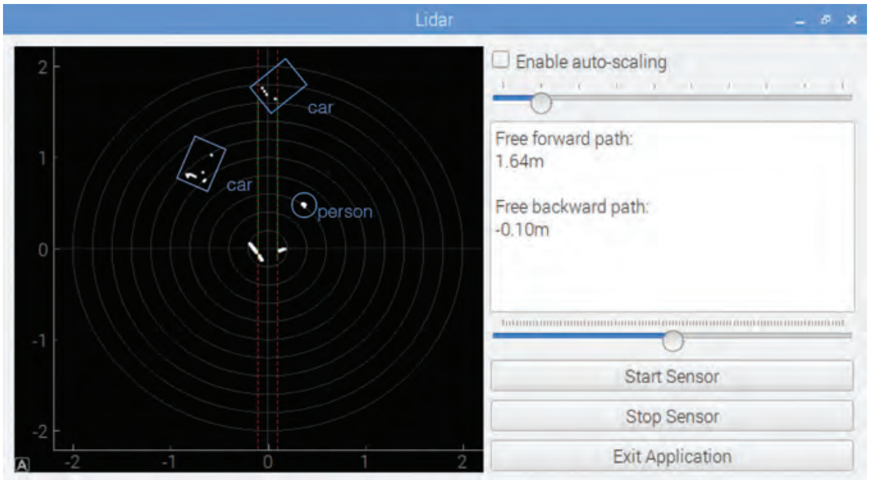

FIGURE 3. A lidar image of two model cars and one doll located in front of the lidar sensor.

Figure 3 depicts the same scenario as detailed in the previous sections, recorded by the rotating lidar scanning sensor. To put the focus of the lidar scan on the described scene, the back of the lidar scanner was covered by a plate. It can be seen that all of the relevant object positions are visible with high accuracy of the angle and distance. However, the lidar sensor does not allow for the estimation of the relative velocity of the detected objects as provided by the radar sensor.

Ultrasound

Equally to the sensors presented in the previous two sections, the principle of an ultrasound sensor is based on radiating a signal and evaluating its echoes. Therefore, an ultrasound sensor transmits an ultrasonic wave, which is reflected by the surrounding objects and received by an ultrasound microphone. The time difference between transmitting and receiving the signal is used to measure the distance of the object that caused the reflection. In automotive applications, ultrasound sensors are often employed to sense the local area of the vehicle. Hence, their main application is range estimation for parking assistance systems.

In addition to the range information, 3D ultrasound sensors also offer the possibility of estimating the angle of objects. This is achieved by employing a microphone array, which measures the reflected ultrasound signal at different positions simultaneously. To adjust the angular resolution of such a sensor, the number of microphones and the geometry of the microphone array can be adapted. The big advantage of ultrasound sensors is their low price, while the main disadvantage is the highly limited maximum range. In Figure 4, a potential scan of the nearby surroundings in a parking maneuver recorded by multiple ultrasound sensors distributed all over the car is shown.

Existing hardware platforms

As already mentioned, the Miniaturized ADAS platform has been integrated into multiple vehicles, each implying its own challenges and capabilities. Compiling courses based on a train allows us to limit the complexity of the individual tasks. Choosing a high-end model car, on the other hand, will enable contributing students to develop a self-driving car for a miniaturized environment. In general, all of the features of the hardware platforms that will be described in the following are fully implemented. Depending on the learning objectives of the practical course, specific software modules are removed and need to be implemented by the students themselves. To provide a better impression on how such a practical course could potentially be realized, we will present chosen hardware platforms as well as examples of their intended functionality.

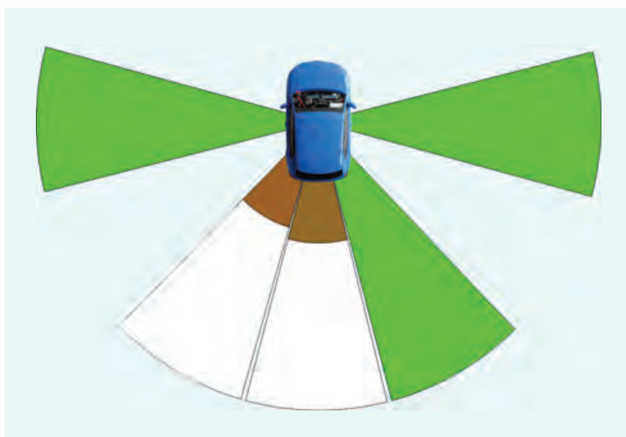

FIGURE 4. A visualized ultrasound measurement during a parking maneuver performed by the Audi Q2 model car.

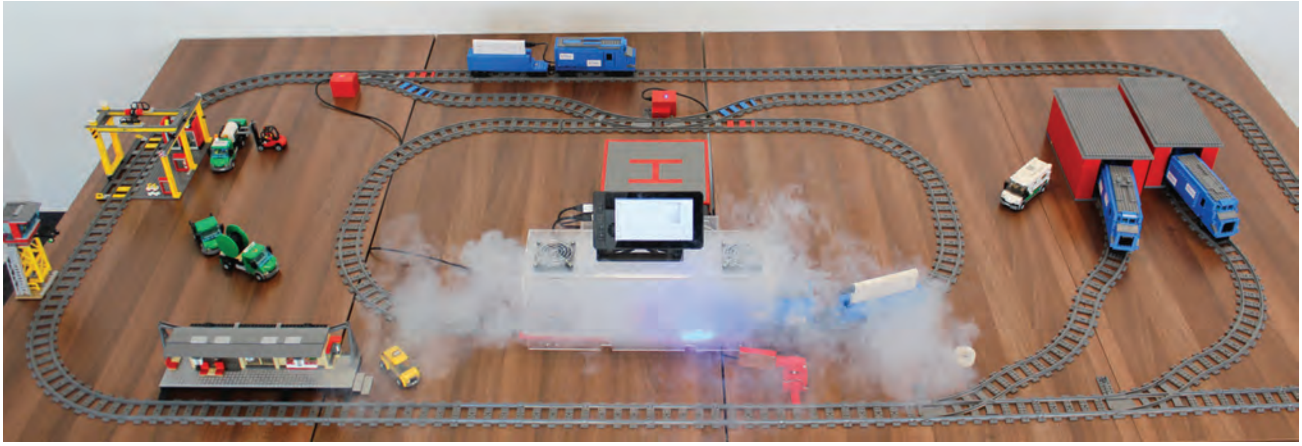

FIGURE 5. The LEGO radar train platform enabling multiple trains on the track, wireless charging, and emergency braking in a foggy environment.

LEGO radar train

The LEGO radar train represents a point of access into Miniaturized ADAS. The reason why we chose a LEGO train for our fundamental application is that trains are bound to rails. This means that there is no necessity to worry about steering, which automatically reduces the problem definition of automated driving by one dimension. Furthermore, LEGO is widely used, perfectly suited for building self-designed model vehicles, and, therefore, also encourages creativity. The first version of the train was developed in cooperation with two high school students from HTL Wels in the course of their final examination project in 2018. Their enthusiasm led to a workshop on radar-based ADASs for HTL students, which, since then, has been held twice at HTL Leonding, with more than 60 participating students each [11]. Furthermore, the same hardware platform in combination with different work packages will be used for an introductory course for the bachelor's degree program in electronics and information technology at JKU.

Meanwhile, the first demonstrator was further evolved, and the latest version offers several features like wireless charging and the associated automatic exchange of two trains on the track. As a highlight, this advanced version is using its integrated 60-GHz radar sensor to perform an emergency brake in a foggy environment, which is illustrated in Figure 5. This scenario clearly highlights the benefits of radar sensors over other sensors like cameras or lidar.

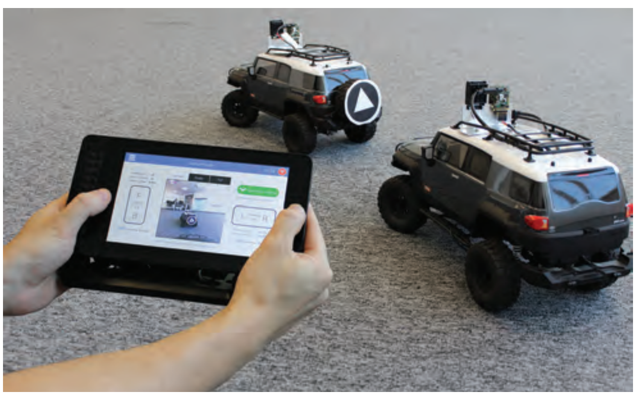

FIGURE 6. The radar model car following another car at a certain distance utilizing a 77-GHz automotive radar sensor and a wide-angle camera.

Radar model car

The next expansion stage of Miniaturized ADAS is a 1:10 model car, initially developed with the University of Applied Sciences Upper Austria Campus Hagenberg in the course of a student project. This practical course is an obligatory part of the curriculum and is often supervised by external lecturers from industry. Within this course, a group of six students attending the bachelor's degree program Hardware–Software–Design automated a commercially available RC model car under the supervision of an Infineon employee. The model car, depicted in Figure 6, also utilizes the described hardware platform based on a Raspberry Pi.

In contrast to the LEGO radar train, the radar model car uses a 77-GHz automotive radar sensor in combination with a wide-angle camera. Basically, it is possible to control the car via a touch screen, which also enables the user to activate and deactivate its assistance systems, such as AEB and ACC. Ultimately, by performing sensor fusion of radar and camera data, the car is able to follow another car at a certain distance. Currently, the steering in this mode is based on the tracking of a predefined symbol utilizing conventional image processing methods.

Beyond that, we supervised a master's thesis on camera-based object detection using CNNs on the Raspberry Pi in combination with an Intel Neural Compute Stick 2 [20]. Due to the promising results, we plan to integrate the investigated machine learning algorithms into the decision logic of the car in one of the next steps.

Currently, we are under discussions with institutes of multiple universities to offer guided practical university courses with a focus on radar signal processing, radar-based ADASs, and sensor fusion strategies. These practical courses will be based on the assignment catalog, which is detailed in the "Course Assignments and Learning Objectives" section.

Audi Q2 model car

Due to the fact that the Raspberry Pi provides only limited processing power, we decided to also include the previously mentioned Audi Q2 into the project. Surely, due to its high price, it naturally conflicts with our presented low-cost platform.

Nevertheless, this model car represents a high-end model car that is well-suited for autonomous driving in a miniaturized environment. The significant computational power available in the Q2 as well as the attached sensors allow for the implementation of highly sophisticated signal processing algorithms.

In addition, the car enables the use of more powerful neural networks for object detection and tracking compared to the radar model car. To push the limits even further, we enhanced the Audi Q2 with a 77-GHz automotive radar sensor in the course of a master's thesis in cooperation with the University of Applied Sciences Upper Austria Campus Hagenberg (see Figure 7).

Course assignments and learning objectives

Based on the three demonstrators, we offer an assignment catalog reaching from imparting the fundamental basics of the individual sensors all the way to the development of highly sophisticated techniques on the system level. Depending on the degree of education and the learning objectives of the course, lecturers are able to arbitrarily choose and modify the desired assignments. Due to the fact that the underlying demonstrators are fully operational, it is possible to provide individual features that are required for specific tasks on higher levels as black-box objects. This enables the students to focus on the given tasks while simultaneously experiencing a comparatively easy and fast sense of achievement. In addition, we encourage the course instructors to use our provided sample solutions as a basis for a detailed discussion at the end of each assignment.

In the following, we describe selected modules from our assignment catalog that are intended to be used for practical university courses on radar-based ADASs. While the first four modules are described in a certain degree of detail, we have decided to give only a brief outline on the remaining modules as detailed explanations would otherwise go beyond the scope of this article. Furthermore, for courses with a higher focus on the system level, this catalog may be extended by modules covering the implementation of an ACC or the investigation of sensor fusion strategies by implementing a radar- and camera-based follow-me scenario.

Module 0: An introduction to the Miniaturized ADAS software framework

The objective of this introductory module is to show students how to remotely control the Raspberry Pi and program the used hardware platform. To introduce the students to the software framework and familiarize them with the provided high-level functions, the main part of this assignment is to execute short test programs and experience their behavior directly on the vehicle.

Module 1: Radar basics

This task starts with a short introduction to FMCW radar systems, explaining the basic functionality and the principles of estimating the range, relative velocity, and angle of surrounding objects. Therefore, the Tx signal is modulated as linear chirp, which can be described by

$$s_{\mathrm{Tx}}(t) = A_{\mathrm{Tx}} \cdot \cos\left(2\pi f_0 t + \pi k t^2\right), \tag{1}$$

for $t \in [0, T_{\mathrm{CH}}]$, where T_{CH} represents the time duration of one chirp period. Furthermore, A_{Tx} represents the amplitude, f_0 the starting frequency of the Tx signal, and $k = B/T_{\mathrm{CH}}$ is the frequency slope of the chirp, whose frequency sweep is denoted as B.

The Tx signal is emitted by a Tx antenna and reflected by the objects in the channel. The superposition of all of these object reflections is sensed by the Rx antennas and further multiplied with the instantaneous Tx signal in a mixer. This finally leads to the intermediate frequency (IF) signal, which contains the required information. To slowly introduce the students to the topic and allow them to investigate the impact of the individual ramp parameters, we provide a small demo program. This program demonstrates how the processed data can be visualized to manually extract the required information from the plots.

Module 2: Range estimation

In this module, the students should estimate the distance of an object that is placed in front of the sensor. Therefore, the basics of spectral estimation techniques, especially the FFT as well as some basic peak finding algorithms, are explained. These methods need to be applied to the IF signal, which, for a single object, can ideally be written as

$$s_{\mathrm{IF}}(t) = A_{\mathrm{IF}} \cdot \cos\left(2\pi f_B t + \Phi\right), \tag{2}$$

FIGURE 7. An Audi Q2 model car executing a lane detection algorithm.

where A_{IF} and Φ are the amplitude and a constant phase term, respectively. In addition, f_B represents the beat frequency, which is directly proportional to the (one-way) distance d between the sensor and the object according to

$$f_B = k \cdot \frac{2d}{c_0}, \tag{3}$$

with c_0 denoting the speed of light.

Module 3: Range Doppler processing

This assignment deals with moving objects and teaches students how the resulting Doppler frequency can be determined by transmitting multiple consecutive ramps. For a dynamic scenario, the IF signal for a single ramp is given by

$$s_{\mathrm{IF}}(t) = A_{\mathrm{IF}} \cdot \cos\left(2\pi(f_B + f_D)t + \Phi\right), \tag{4}$$

containing the Doppler frequency,

$$f_D = f_0 \frac{2v}{c_0}, \tag{5}$$

with v describing the relative velocity of the moving object. The 2D data set resulting from transmitting several ramps needs to be processed via a 2D FFT to obtain a range-Doppler map.

Module 4: Object detection

Due to the fact that a constant threshold to separate actual objects from noise is not sufficient in a constantly changing environment, the main part of this assignment is represented by CFAR algorithms. After an introduction of the concept, the students should implement an ordered statistic CFAR algorithm and test if the vehicle is able to automatically detect the object and estimate its distance and velocity.

Module 5: Angle estimation

After an object is detected in the range-Doppler map, the 2D FFT needs to be applied to the data of all of the remaining Rx channels. By performing another FFT across the third dimension of this 3D matrix, at the peaks detected by the CFAR algorithm, the angle of incidence of the reflected waves can be estimated.

Module 6: Clustering algorithms

Due to the fact that complex-shaped objects typically produce several reflections, it is essential to figure out which detections belong to the same object. In this module, the students learn about density-based spatial clustering of applications with noise to group all of the detections into potential clusters and generate an object list. Finally, the actual positions of the objects can be visualized, as illustrated in Figure 2.

Module 7: Tracking algorithms

As the final assignment on sensor level, the goal of this module is to implement a Kalman filter, which is used to increase the confidence in a potential detection. This filter is used to eliminate false detections that solely accrue in a single measurement frame.

Module 8: Radar-based emergency braking

In this system-level assignment, the students need to use the postprocessed radar data to appropriately control the motor of the vehicle. In the first step, the vehicle should perform an emergency brake if an object is detected at a certain distance. In the second step, the emergency braking should be triggered only if the obstacle is directly in front of the sensor and avoid unjustified braking if the obstacle is located on a different lane or railroad track.

Achievements, feedback, and lessons learned

Utilizing the LEGO radar train, we established a workshop on radar-based ADASs for HTL students, which, since then, has been held twice, with more than 60 participating students each [11]. In Figure 8, we give insight into one of these workshops to convey the atmosphere. However, throughout the first workshop, we identified a few weaknesses in our platform that could successfully be eliminated, resulting in a further improvement of Miniaturized ADAS. In the following, we briefly share our experience.

To keep the individual groups as small as possible, we provided 10 workshop kits based on the Raspberry Pi 3, which is solely able to wirelessly communicate via a 2.4-GHz Wi-Fi connection. Accordingly, we faced the issue that this frequency band was partially overloaded, and thus, communication problems occurred. To overcome this problem, we upgraded our workshop kits to the Raspberry Pi 3B+, which also supports the 5-GHz frequency band, for the second workshop.

In addition, we also experienced that the students had problems in efficiently utilizing the four cores of the ARM processor available on the Raspberry Pi. This is required to enable real-time capability and process the data of multiple sensors simultaneously. Therefore, we developed a co-routine-based software framework in combination with message passing systems [11]. This implementation allows for the design of concurrent systems, avoiding the challenges usually associated with traditional multithreading or multiprocessing techniques, and consequently offers an easy-to-use experience.

FIGURE 8. The 2019 workshop at HTL Leonding had more than 60 participating students. (Source: [11].)

Despite the described problems that occurred during the first workshop, we received consistently positive feedback from the contributing students. They especially liked the fact that they were able to directly test their implemented solutions on the provided hardware. Furthermore, most of the contributing students were familiar with the ADAS in modern cars and were really excited to reproduce such systems in a miniaturized version. However, in our opinion, the best feedback was the ambition, commitment, and creativity with which the students solved the given tasks.

The opportunity to exhibit and present their implemented sample solutions as demonstrators at various trade shows in Austria and Germany was another reward that was highly appreciated by our students who contributed to the development of Miniaturized ADAS platform.

Outlook and future work

Currently, we are supervising a bachelor's thesis in cooperation with the University of Applied Sciences Upper Austria Campus Hagenberg with the objective of open sourcing the LEGO radar train workshop on GitHub [30]. This will include the part lists, building instructions, schematics, printed circuit board designs, corresponding exercises, and a reference implementation.

In addition, we revised the workshop to be held as an introductory course for the bachelor's degree program in electronics and information technology at JKU, which will be held for the first time during the winter semester of 2021/2022.

As a further step, our objective is to prepare guided practical university courses based on the assignment catalog presented in the "Course Assignments and Learning Objectives" section, with a particular focus on radar signal processing, radar-based ADASs, and sensor fusion strategies, which will be realized on the radar model car. Once this practical university course is successfully introduced at JKU, it shall also be open sourced on GitHub [30].

Unfortunately, due to the COVID-19 pandemic, we were not able to hold the workshop in 2020. Therefore, we adapted the workshop such that the students are able to remotely program the train from home and observe their progress on the actual hardware via a camera stream. This setup has already successfully been tested at a virtual event called "Industry Meets Makers," and we plan to use it for the upcoming HTL workshops in 2021.

Conclusions

In this article, we presented an innovative educational platform for ADASs and autonomous driving for a miniaturized environment with the aim of inspiring other academic institutions around the world to either collaborate on our platform or develop their own to adequately support students in their education. We further provided an overview of the most promising sensors currently used to automate a vehicle and thereby demonstrated the potential fields of applications for signal processing and machine learning in this area. To give an impression on how Miniaturized ADAS can be used for teaching, we finally presented an exemplary assignment catalog for radar-based ADASs as well as the provided hardware platforms.

Acknowledgments

We would like to thank Infineon Technologies for supporting this work and allowing for its publication. We also express our sincere thanks to Alexander Reisenzahn, Stefan Matzinger, Gerhard Riess, and Manfred Ruhmer, who have always actively supported the Miniaturized ADAS project. Finally, we would like to thank all of the contributing students for being part of this project.

Authors

Michael Gerstmair (michael.gerstmair@infineon.com) received his Dipl.-Ing. in information electronics from Johannes Kepler University (JKU) Linz, Austria, in 2016. Currently, he works as a radar system engineer at Infineon Technologies Linz, 4040, Austria. From 2016 to 2020, he was a Ph.D. student at Infineon Technologies Linz in cooperation with the Institute of Signal Processing at JKU, Austria. His research focuses on radio-frequency impairments in highly integrated radar front ends as well as radar signal processing for automotive applications. He has supervised various bachelor's and master's theses related to radar- and camera-based advanced driver assistance systems and worked as a part-time lecturer at the University of Applied Sciences in Hagenberg, in 2018. He is a Student Member of IEEE.

Martin Gschwandtner (martin.gschwandtner@infineon.com) received his master's degree in hardware/software systems engineering in 2004 and in information engineering and management in 2009, both from the University of Applied Sciences in Hagenberg, Austria. Since 2004, he has been with Infineon Technologies Linz, 4040, Austria. From 2017 to 2019, he was the head of automotive radar software development. He is currently responsible for R&D funding and innovation activities. He has been involved in student activities for more than 10 years and has mentored in cooperation with technical colleges and universities more than 20 thesis and student projects in his career.

Rainer Findenig (rainer.findenig@infineon.com) received his Ph.D. degree from Johannes Kepler University Linz, Austria, in 2016 with Austria's highest possible honor in education, the Promotio sub auspiciis Praesidentis rei publicae. Currently, he is with Infineon Technologies Linz, 4040, Austria, where he is driving the next-generation automotive radar sensors' control architecture. Additionally, he has more than a decade of experience in teaching system modeling and software and hardware development at the Upper Austrian University of Applied Sciences, initially in a full-time research and teaching position and now as an external lecturer. He has supervised more than a dozen bachelor's and master's theses and several student projects.

Oliver Lang (oliver.lang@jku.at) received his Ph.D. degree from Johannes Kepler University (JKU) Linz, Austria, in 2018. He received his bachelor's degree in electrical engineering and information technology and his master's degree in

microelectronics from the Vienna University of Technology, Austria, in 2011 and 2014, respectively. Since 2019, he has been a university assistant at the Institute of Signal Processing at JKU, Linz, 4040, Austria, where he holds courses in the field of digital and statistical signal processing. He is the main inventor of several patents and patent applications in the field of automotive radar systems. He is a Member of IEEE.

Alexander Melzer (alexander.melzer@infineon.com) received his Dipl.-Ing. degree from the Graz University of Technology in 2012. From 2013 to 2014, he was with Maxim Integrated Austria. From 2014 to 2017, he worked on his Ph.D. thesis in the field of radar signal processing at Johannes Kepler University Linz in cooperation with Infineon Technologies. During his Ph.D. thesis, he published several conference and journal papers and holds various patents in the area of phase noise estimation and leakage cancelation in radar transceivers. Among various awards, he received the German Society of Information Technology Award in 2016. Since 2017, he has been a product architect for future automotive radar transceiver monolithic microwave integrated circuits at Infineon Technologies Linz, 4040, Austria.

Mario Huemer (mario.huemer@jku.at) received his Dipl.-Ing. degree in mechatronics and his Dr.techn. degree from Johannes Kepler University (JKU) Linz, Austria, in 1996 and 1999, respectively. After holding positions in industry and academia, he became an associate professor at the University of Erlangen-Nuremberg, Germany, from 2004 to 2007, and a full professor at Klagenfurt University, Austria, from 2007 to 2013. Since September 2013, he has been a full professor and head of the Institute of Signal Processing at JKU Linz, 4040, Austria. His research interests include statistical and adaptive signal processing with applications in information and communications engineering, in particular radio-frequency transceivers for communications and radar, sensor, and biomedical signal processing. He is a Senior Member of IEEE.

References

[1] W. J. Gordon, *Our Home Railways, How They Began and How They Are Worked*, vol. 1, London and New York: Fredrick Warne, 1910.

[2] T. Swan. "Ford's assembly line turns 100: How it really put the world on wheels." Apr. 2013. Car and driver. https://www.caranddriver.com/features/a15115930/fords-assembly-line-turns-100-how-it-really-put-the-world-on-wheels-feature/ (accessed Jan. 17, 2020).

[3] M. Gerstmair, A. Melzer, A. Onic, and M. Huemer, "On the safe road toward autonomous driving: Phase noise monitoring in radar sensors for functional safety compliance," *IEEE Signal Process. Mag.,* vol. 36, no. 5, pp. 60–70, Sept. 2019. doi: 10.1109/MSP.2019.2902914.

[4] "2018 road safety statistics: What is behind the figures?" Apr. 2019. European Commission. https://ec.europa.eu/commission/presscorner/detail/en/MEMO_19_1990 (accessed Jan. 17, 2020).

[5] D. Kolb, *Experiential Learning: Experience as the source of Learning and Development,* 2nd ed., New York: Pearson, Jan. 2015.

[6] M. Nolte, T. Form, S. Ernst, R. Graubohm, and M. Maurer, "The Carolo-Cup student competition: Involving students with automated driving," in *Proc. 2018 12th Euro. Workshop Microelectron. Educ. (EWME)*, Sept. 2018, pp. 95–99. doi: 10.1109/EWME.2018.8629462.

[7] The NXP-cup EMEA. NXP Community. https://community.nxp.com/groups/tfc-emea (accessed Jan. 14, 2020).

[8] "Audi autonomous driving cup 2018." 2018. Audi AG. https://www.audi-autonomous-driving-cup.com (accessed Jan. 14, 2020).

[9] "Carolo-Cup regulations 2019," 2018. TU Braunschweig, https://www.tu-braunschweig.de/fileadmin/Redaktionsgruppen/Institute_Fakultaet_5/Carolo-Cup/180710_Carolo-Cup_Regulations.pdf

[10] Wikipedia contributors. Höhere Technische Lehranstalt." Wikipedia. [Online]. Available: https://en.wikipedia.org/wiki/Hoehere_Technische_Lehranstalt (accessed Jan. 17, 2020).

[11] M. Gerstmair, M. Gschwandtner, R. Findenig, A. Melzer, and M. Huemer, "LEGO radar train–An educational workshop on radar-based advanced driver assistance systems," in *Proc. 28th Euro. Signal Process. Conf. (EUSIPCO),* pp. 1981–1985. doi: 10.23919/Eusipco47968.2020.9287353.

[12] A. K. Jain, "Working model of self-driving car using convolutional neural network, raspberry Pi and Arduino," in *Proc. 2nd Int. Conf. Electron., Commun. Aerosp. Technol. (ICECA),* Mar. 2018, pp. 1630–1635. doi: 10.1109/ICECA.2018.8474620.

[13] B. Blaga, M. Deac, R. W. Y. Al-doori, M. Negru, and R. Dǎnescu, "Miniature autonomous vehicle development on raspberry pi," in *Proc. IEEE 14th Int. Conf. Intell. Comput. Commun. Process. (ICCP),* Sept. 2018, pp. 229–236. doi: 10.1109/ICCP.2018.8516589.

[14] "Volunteers of open source project." Donkey car. https://www.donkeycar.com (accessed Jan. 14, 2020).

[15] N. Baras, G. Nantzios, D. Ziouzios, and M. Dasygenis, "Autonomous obstacle avoidance vehicle using lidar and an embedded system," in *Proc. 8th Int. Conf. Modern Circuits Syst. Technol. (MOCAST),* May 2019, pp. 1–4. doi: 10.1109/MOCAST.2019.8742065.

[16] S. Zug, C. Steup, J. Scholle, C. Berger, O. Landsiedel, F. Schuldt, J. Rieken, R. Matthaei et al., "Technical evaluation of the Carolo-Cup 2014 - A competition for self-driving miniature cars," in *Proc. IEEE Int. Symp. Robot. Sens. Environ. (ROSE),* Oct. 2014, pp. 100–105. doi: 10.1109/ROSE.2014.6952991.

[17] R. Adomat, G.-O. Geduld, M. Schamberger, and P. Rieth, "Advanced driver assistance systems for increased comfort and safety–Current developments and an outlook to the future on the road," in *Advanced Microsystems for Automotive Applications 2003,* J. Valldorf and W. Gessner, Eds. Berlin: Springer-Verlag, 2003, pp. 431–446.

[18] F. Bounini, D. Gingras, V. Lapointe, and H. Pollart, "Autonomous vehicle and real time road lanes detection and tracking," in *Proc. IEEE Veh. Power Propulsion Conf. (VPPC),* 2015, pp. 1–6. doi: 10.1109/VPPC.2015.7352903.

[19] R. C. Gonzalez, "Deep convolutional neural networks [Lecture Notes]," *IEEE Signal Process. Mag.,* vol. 35, no. 6, pp. 79–87, 2018. doi: 10.1109/MSP.2018.2842646.

[20] S. Rader, "Lightweight convolutional neural networks for camera-based object detection on mobile devices," Master's thesis, Johannes Kepler Univ., 2020.

[21] M. Gerstmair, A. Melzer, A. Onic, R. Stuhlberger, and M. Huemer, "Highly efficient environment for FMCW radar phase noise simulations in if domain," *IEEE Trans. Circuits Syst. II, Express Briefs,* vol. 65, no. 5, pp. 582–586, 2018. doi: 10.1109/TCSII.2018.2821270.

[22] S. M. Patole, M. Torlak, D. Wang, and M. Ali, "Automotive radars: A review of signal processing techniques," *IEEE Signal Process. Mag.,* vol. 34, no. 2, pp. 22–35, Mar. 2017. doi: 10.1109/MSP.2016.2628914.

[23] H. Rohling, "Ordered statistic CFAR technique–An overview," in *Proc. 12th Int. Radar Symp. (IRS),* 2011, pp. 631–638.

[24] J. Dickmann, N. Appenrodt, J. Klappstein, H. L. Bloecher, M. Muntzinger, A. Sailer, M. Hahn, and C. Brenk, "Making Bertha see even more: Radar contribution," *IEEE Access,* vol. 3, pp. 1233–1247, Aug. 2015. doi: 10.1109/ACCESS.2015.2454533.

[25] J. Hasch, E. Topak, R. Schnabel, T. Zwick, R. Weigel, and C. Waldschmidt, "Millimeter-wave technology for automotive radar sensors in the 77 GHz frequency band," *IEEE Trans. Microw. Theory Techn.,* vol. 60, no. 3, pp. 845–860, 2012. doi: 10.1109/TMTT.2011.2178427.

[26] R. Schnabel, D. Mittelstrab, T. Binzer, C. Waldschmidt, and R. Weigel, "Reflection, refraction, and self-jamming," *IEEE Microw. Mag.,* vol. 13, no. 3, pp. 107–117, Apr. 2012. doi: 10.1109/MMM.2012.2186095.

[27] J. Lien, N. Gillian, M. E. Karagozler, P. Amihood, C. Schwesig, E. Olson, H. Raja, and I. Poupyrev, "Soli: Ubiquitous gesture sensing with millimeter wave radar," *ACM Trans. Graph*, vol. 35, no. 4, July 2016. doi: 10.1145/2897824.2925953.

[28] M. Kutila, P. Pyykönen, W. Ritter, O. Sawade, and B. Schäufele, "Automotive LIDAR sensor development scenarios for harsh weather conditions," in *Proc. IEEE 19th Int. Conf. Intell. Transp. Syst. (ITSC),* 2016, pp. 265–270. doi: 10.1109/ITSC.2016.7795565.

[29] R. Whyte, L. Streeter, M. Cree, and A. Dorrington, "Application of lidar techniques to time-of-flight range imaging," *Appl. Opt.,* vol. 54, no. 33, pp. 9654–9664, Nov. 2015. doi: 10.1364/AO.54.009654.

[30] "Miniaturized ADAS." GitHub. https://github.com/Infineon/miniaturized-adas (accessed Dec. 4, 2020).

SP

Wendy Tang and Petar M. Djuric

Bachelor of Science in Electrical Engineering Online

A journey of challenges and triumphs

©SHUTTERSTOCK.COM/LUCKYSTEP

In June 2005, three State University of New York (SUNY) campuses, namely Stony Brook University, the University at Buffalo, and Binghamton University, joined forces to provide a completely online electrical engineering program that would lead to a bachelor of science in electrical engineering (BSEE) degree. The program was put together to serve a nontraditional student body that is otherwise unable to achieve higher education in electrical engineering. The project was jump-started by an award from the Alfred P. Sloan Foundation. This multi-institutional approach was an effective medium for sharing the resources and revenue of the online offering of courses. The Sloan Foundation grant initiated a big step toward the "Learning Anywhere, Anytime" vision for science, technology, engineering, and mathematics (STEM) education. It helped to set up the infrastructure of an ambitious joint venture among the three SUNY campuses to collaborate on a jointly offered program.

In this article, we chronicle the journey of the program going through two Accreditation Board for Engineering and Technology (ABET) accreditation processes, describe the program's nuances and its differences from massive online courses (MOOCs), explain how it promotes innovation, and provide some details of the offered signal processing courses. We also show how online learning can facilitate and strengthen cross-disciplinary communities of research and education via online collaboration. Further, we argue that it enables the "learning to learn" paradigm, where the teacher also becomes the learner.

In the article and based on our experience, we convey our belief that the online learning paradigm is conducive to 21st-century education, where creativity, critical thinking, communication, and collaboration are encouraged. Our experience thus far has been very rewarding. Some of our colleagues with careers in engineering education of more than 20 years maintain that they have seen the best students in these online courses. Feedback from students has also been very positive, with most of them enjoying the vigor and flexibility of the schedule of our courses.

Digital Object Identifier 10.1109/MSP.2021.3057048
Date of current version: 28 April 2021

While the white paper for this article was prepared, we were still living in pre-COVID-19 times. We had no idea that the academic world would quickly turn upside down as did more or less every aspect of our daily lives. As a result of the pandemic that followed, we have witnessed forced changes in course delivery in academia across the world. On very short notice, instructors had to prepare lectures for online delivery, and, in courses that had laboratory components, they had to come up with creative solutions quickly. There is no doubt that many readers of this article who work in academia have already had the experience of teaching online courses and have drawn their own conclusions on the advantages and disadvantages of this mode of teaching.

A brief description of the program

The "Learning Anywhere, Anytime" vision of education has been widely supported and adopted [1]. This vision has been promoted by various institutions of the federal government of the United States [including the National Science Foundation (NSF) and the U.S. Department of Education through the Fund for the Improvement of Postsecondary Education (FIPSE)], private foundations (such as the Alfred P. Sloan Foundation), and industry (such as the Microsoft Corporation) [2]–[4]. It is also gaining wide acceptance as the U.S. Department of Education states that online learning is one of the fastest-growing trends in educational uses of technology. One of its reports from 2009 concluded that "students who took all or part of their class online performed better, on average, than those taking the same course through traditional, face-to-face instruction" [5].

Until recently, however, most fully online offerings of courses have not been in the STEM disciplines. But with the help from the FIPSE, NSF, and Sloan Foundation, this has changed. The growth of technologies for course management and virtual laboratories as well as software for the remote access of laboratory equipment [6]–[11] has also contributed to the increased number of online offerings of courses in the STEM disciplines in recent years. Nevertheless, most online STEM courses have either been at the graduate level or are "blended" (combining online and face-to-face meetings). The rest of the world was also quickly getting traction in online courses, even in areas with technological barriers.

Some of the driving forces for this trend have been population growth and the exploding demand for education [12]. However, the realization of the vision of "Learning Anywhere, Anytime" for a complete bachelor's degree online in STEM disciplines was still a work in progress until 2020. Then the pandemic came, and all of the disciplines, including STEM, had no choice but to deliver online courses. It remains to be seen how these forced changes will reverberate once academia returns to normal operations.

The project started first as an offering of courses for nonmatriculated students. In May 2011, it switched into an online program leading to a BSEE degree. At that time, Stony Brook University also received approval for the program from the New York State Education Department. A key feature of the upper-division program is that students can complete all of the electrical engineering coursework online without setting foot on campus. This is one of the first completely online electrical engineering degree programs at the bachelor's level in the nation. The program is ABET accredited.

A key feature of the upper-division program is that students can complete all of the electrical engineering coursework online without setting foot on campus.

All of the courses are delivered asynchronously, including the electronic laboratory courses. During the semester, the instructors and the teaching assistants are in close communication with the students in the class, often via scheduled video meetings. The students in the courses with laboratories perform the same experiments as those in traditional offerings (more on this in the "Laboratory-Based Courses" section). This is made possible via the innovative use of technology. Finally, some of our online lectures developed in the program are also being used to supplement and enrich on-campus offerings, hence creating a win–win setting. Students take the basic mathematics, science, and general education courses elsewhere (online or at local institutions). All of these courses are carefully assessed before they are approved for the BSEE degree.

The educational objectives of the program are:

- PEO1: Our graduates should excel in engineering positions in industry and other organizations that emphasize the design and implementation of engineering systems and devices.
- PEO2: Our graduates should excel in the best graduate schools, reaching advanced degrees in engineering and related disciplines.

The expected student learning outcomes are:

1) an ability to identify, formulate, and solve engineering problems by applying the principles of engineering, science, and mathematics
2) an ability to apply engineering design to produce solutions that meet specified needs with consideration of public health, safety, and welfare as well as global, cultural, social, environmental, and economic factors
3) an ability to communicate effectively with a range of audiences
4) an ability to recognize ethical and professional responsibilities in engineering situations and make informed judgments, which must consider the impact of engineering solutions in global, economic, environmental, and societal contexts
5) an ability to function effectively on a team whose members together provide leadership, create a collaborative and inclusive environment, establish goals, plan tasks, and meet objectives
6) an ability to develop and conduct appropriate experimentation, analyze and interpret data, and use engineering judgment to draw a conclusion

7) an ability to acquire and apply new knowledge as needed, using appropriate learning strategies.

In terms of course offering, the program curriculum consists of 17 core courses and four technical electives from a list of eight technical electives. A more detailed description of the courses and their syllabi can be found on the curriculum page of the program's website [13]. This list evolves with time to reflect the changes in the discipline, needs of industry, and interest of the students. Currently, the core courses in the program are as follows:

- EEO 124: C or C++ Programming, 3 credits
- EEO 218: Digital Logic Design, 3 credits
- EEO 219: Digital Logic Design Lab, 1 credit
- EEO 271: Electric Circuits Analysis, 3 credits
- EEO 224: Object Oriented Programming for Electrical and Computer Engineers, 3 credits
- EEO 300: Technical Communications for Electrical Engineers, 3 credits
- EEO 301: Signals and Systems, 3 credits
- EEO 302: Engineering Ethics and Societal Impact, 3 credits
- EEO 306: Random Signals, 3 credits
- EEO 311: Electronics II, 3 credits
- EEO 315: Electronics I, 3 credits
- EEO 323: Electromagnetics, 3 credits
- EEO 331: Semiconductor Devices, 3 credits
- EEO 352: Electronics Lab I, 3 credits
- EEO 353: Electronics Lab II, 3 credits
- EEO 440: Engineering Design I, 3 credits
- EEO 441: Engineering Design II, 3 credits.

The technical electives (four courses are to be selected from the list below) are as follows:

- EEO 304: Electronic Instrumentation and Operational Amplifiers, 3 credits
- EEO 303: Digital Signal Processing, 3 credits
- EEO 314: MOS Transistor Modeling, 3 credits
- EEO 316: Integrated Electronic Devices and Circuits, 3 credits
- EEO 346: Computer Communications, 3 credits
- EEO 366: Embedded Mixed-Signal Systems, 3 credits
- EEO 388: Foundations of Machine Learning, 3 credits
- EEO 414: Fundamentals of Low Noise Electronics for Sensors
- EEO 425: Electric Machinery and Energy Conversion, 3 credits
- EEO 470: Renewable Distributed Generation and Storage, 3 credits
- EEO 482: Power Systems Engineering, 3 credits.

These courses are offered via the Blackboard Learn environment. Some courses have PowerPoint slides with voice and annotations, some use Echo 360 software to capture on-campus lectures with PowerPoint slides, while others use different technologies [14].

Most of the students in the BSEE program are nontraditional students. They are working professionals pursuing the degree part time for career advancement while supporting families. In the fall of 2020, we had 37 students, 32 of them part-time students and five of them full-time students. Of the 32 part-time students, most were engineering technicians pursuing the degree to become engineers. We also had a couple of female students working toward the degree to return to the workforce after raising a family as well as one patent lawyer having as a goal the degree to expand his practice to include electrical engineering.

Most of the students in the BSEE program are nontraditional students.

Our prospective students are primarily working professionals. However, in the fall of 2020, due to the pandemic, we also had a few traditional students who had completed an associate of science degree from a community college. In principle, the program accepts other students, too, as long as they satisfy our admission requirements.

ABET evaluations of the program

In May 2014, the program graduated its first cohort of three students, and a team of three ABET evaluators visited campus. This was ABET's first accreditation of a completely online undergraduate program in electrical engineering. In addition to Stony Brook University faculty, faculty members from the University of Buffalo and Binghamton University were onsite to meet the accreditation team. In October 2017, the program went through its second ABET accreditation visit along with the onsite engineering programs of the university. The program passed the accreditation visit with flying colors.

In our first ABET accreditation, while the evaluators were impressed with the strength of the program and the dedication of the faculty, there was a shortcoming in that students in the program did not have enough exposure to multidisciplinary team experience. We argued that most students in the program were working professionals who had ample team experience at their jobs. The ABET evaluators agreed that the team experience from work was appropriate but that the experience needed to be documented. Since then, the program has introduced an internship course (EEO 488: Internship) in which students receive credits for work related to projects. Part of the requirements of the course is the documentation of team experience.

In our second ABET accreditation, when the program was evaluated with the onsite programs, we found that, by and large, the accreditation requirements were quite similar. However, in the BSEE program, there were no paper assignments. All of the coursework was submitted electronically, which facilitated the assessment process. The evaluators found no shortcomings in the program.

MOOCs and BSEE online program comparisons

MOOCs are online courses free for anyone to enroll in and learn about topics of almost any type. Thus, MOOCs may appear to be in direct competition with programs like ours. We argue otherwise.

MOOCs as a distant education framework was introduced in 2008 [15], and it gained popularity soon after. In 2011, 450,000 students took three computer science courses offered by Stanford University [16], and in 2012, *The New York Times* declared 2012 the "Year of the MOOC" [17]. In the following years, the number of students and the number of offered courses continued to swell at a staggering pace. According to ClassCentral [18], in 2019, there were 110 million students (excluding China) who took online courses from a list of about 13,500 courses. Further, in 2019, providers introduced more than 2,500 courses [18]. The courses are offered by many universities and are delivered via respectable platforms that provide the technology for distributing the video lectures with possible follow-up tests. Some of the courses are interactive, and some offer user forums and social media discussions.

Programs like ours are distinguished from those offered by MOOCs in that they are for undergraduate students and they have laboratory components in several courses.

MOOCs have also introduced online master's degrees that can be completed fully online [19] (in 2019, 11 new online degrees and 170 microcredentials were introduced [18]). There are now more than 50 programs, and they are put together in partnerships with universities. Further, these degrees are accredited and recognized, and their cost is lower than that of other online and on-campus degrees. Popular master of science degree programs are in computer science and engineering, business and management, data science and analytics, cybersecurity, IT management, and public health and health care.

Still, the excitement about MOOCs seems to have waned as can be attested by a slowdown in the increase of new users. On the other hand, MOOC platforms are experiencing a growth in paying customers, which perhaps indicates a shift in the focus of MOOCs.

Programs like ours are distinguished from those offered by MOOCs in that they are for undergraduate students and they have laboratory components in several courses. They are structured in a similar way to regular on-campus programs such that the difference between them is minimized. The assessment and evaluation in the BSEE program are similar to the onsite program. For the most part, evaluations in the courses are based on written answers for traditional engineering questions, including math derivations, circuit diagrams, and designs, as well as submitting written reports for laboratory courses and projects. Presentations in project-oriented courses, such as the capstone design project, are an important component of assessment. All of the exams are timed and proctored.

The admission criteria for the BSEE program are much stricter than those for MOOCs. Although the MOOC platforms have a large number of free electrical engineering undergraduate courses, grades in these courses are usually based on automatically graded assessments or peer-to-peer reviews. The exams in the BSEE program are proctored in a way that would maximize authenticity and are graded by professors or teaching assistants. The students receive feedback from their instructors on a regular basis and can speak to them directly during office hours or during individually arranged video meetings. An important course in the program is the senior design project, which is a two-semester course where students work on teams to solve real-world engineering problems. This course in the BSEE program is designed in such a way that the activities are indeed the result of teamwork.

As already pointed out, the quality of the classes in the online program is similar to that of the on-campus classes, and, thus, the online students in the program have to invest significant time and effort to complete the courses. Further, students in the classes are encouraged to interact with their classmates to share experiences. This is a much more natural process when the number of students is small as opposed to the large groups of students in MOOCs. If we add to this the fact that the backgrounds of the students in the BSEE program are much more congruent than in any MOOCs in the context of meeting the goals of the program, this is very helpful. The students can get much more out of their classes when they are with peers that face similar challenges and have similar backgrounds.

Last but not least, the BSEE program undergoes periodic accreditations that are designed to provide assurance that the program meets the quality standards of the profession. This, in itself, offers the students more respected credentials for the completed degree.

Laboratory-based courses

The online delivery of laboratory-based courses is a challenge. However, in the electrical engineering discipline, laboratory courses often involve the building of low-power electronic circuitry that is powered by dc batteries, which removes safety concerns. Furthermore, technological advances have made USB oscilloscopes available for reasonable costs. The pandemic of 2020 has created a strong demand for equipment that allows students to perform experiments in their homes, and that has pushed many manufacturers to quickly step in and fill the existing void with novel products.

In the BSEE program, we have three required laboratory courses: EEO 219: Digital Electronics Lab, EEO 352: Electronics Lab I, and EEO 353: Electronics Lab II, all of which require students to purchase a commercial laboratory kit that costs about US$300. In these courses, they learn to use the instruments, implement designs, rectify problems, optimize circuits, debug, and perform characterizations. They also learn how to analyze data, evaluate the significance of their findings, and prepare reports and presentations. Since their inception, the goal of the electronics laboratories has been to provide students with hands-on experiences.

Students perform circuit design and simulations using software and then conduct experiments with the hardware. They use computer-based instruments locally most of the time. On rare occasions, e.g., when they perform device characterizations requiring fine-time resolution or high frequencies, they can access equipment online. This combination of experiments improves with every passing year. We build on past experience and constantly improving technology for delivering online education with the aim of providing our students with a rich set of opportunities to learn by doing. More details of our laboratory courses can be found in [14].

The overall aim is to keep the level of education in the online classes at least as high as that in the onsite classes.

Senior design projects are another challenge. Due to the geographical separation, most students opt for an individual capstone design project. As pointed out, during our first ABET accreditation visit in 2014, the evaluators expressed concern about the lack of teamwork. We argued that our students in the BSEE program are working professionals who have already demonstrated teamwork experience in their professional careers. The concern was removed after we introduced an internship course in which students discuss their teamwork experience.

Based on our experience, we found that, in most cases, students do individual projects. However, due to the fact that most of these students are mature professionals, a large number of them proposed their own projects based on their personal and professional experiences. From the past, some interesting design projects included a tankless water heating system; a smart garden system that monitors soil conditions to optimize planting times, watering levels, fertilization schedules, and so on; an innovative digital electromechanical timing device to control the motor runtime for pressurizing a plant; and a programmable real-time audio processor using a commercial field-programmable gate array board. Each student in this course has one-on-one consultation with the project advisor, mostly on a weekly or biweekly basis. It is our impression that, for many of these projects, the design is more complex than that of the on-campus projects. We explain this by the more advanced professional experience of the BSEE students.

We have also explored the option of connecting the BSEE online students with a team from the traditional on-campus program. In these cases, the BSEE online students were often responsible for the software component of the project. We observed interesting dynamics between the BSEE students, who are often much more senior than the traditional onsite students, and the on-campus students, who are typically in their early twenties. Often the BSEE students provide more well-rounded input drawn from their work experience, while the traditional students excel in mathematical and analytical skills. We also have had some teams of two or more BSEE students working together on their capstone projects. However, in this case, the results have been mixed. In general, it seems that the teams that had members who live within relatively close geographic proximity did well as they were able to arrange face-to-face meetings. The teams with members that have been geographically far away from each other have not fared as well.

Signal processing-related courses

The program has two required courses related to signal processing (EEO 301: Signals and Systems and EEO 306: Random Signals) and one elective course (EEO 303: Digital Signal Processing). We provide details of how these courses are offered and how we increase student engagement during the semester and make the engagement continuous.

The syllabus of the course on signals and systems includes all of the standard topics that one would expect in such a course, including an introduction to continuous-time and discrete-time signals and linear systems, differential and difference equations, convolution, Fourier series and transforms, transfer functions, frequency responses, the Laplace and z transforms. In the course on random signals, the emphasis is on introducing the concepts of random experiments and events, random variables, probability distribution and density functions, random processes, and the response of linear systems to random inputs. Some of the basics of statistics are also provided. The elective course on digital signal processing covers ordinary topics.

The courses are well structured, and the lectures are posted once or twice a week (as if the students attend the lectures on campus). The lectures contain hands-on examples that follow the theory and are connected with real-world applications. They are accompanied by homework assignments that may include exercises that need to be carried out using MATLAB or other programming languages. The instructors hold office hours at times that are usually after hours so that students who have full-time jobs can have the opportunity to speak directly to them and ask questions. Instructors also often arrange for meetings that are outside of office hours. Such arrangements are possible because the number of students in a class is not large.

The exams are of the same form as the exams for on-campus students; they are proctored and usually take place over the weekends. The exams are graded by the instructors in the usual way and returned to the students. The pace of the instruction is identical to that of on-campus students, and it occurs over a period of 14 weeks. The overall aim is to keep the level of education in the online classes at least as high as that in the onsite classes.

Challenges and triumphs

In our last 10 years of experience delivering online courses, we have encountered a number of challenges. They include authentication/verification, student engagement, and student/faculty interactions. At the same time, we have also gained insights into some unexpected benefits of the online offerings.

The authentication and verification of student identity are vital to the integrity of any program. In the early stages of the program, we came across students who engaged in academic dishonesty when an unqualified person was administering an exam. This was promptly changed to a proctoring policy where every exam, including midterms and tests, must be proctored by a competent delegate. The person must be the faculty/staff of another higher education institution or a full-time employee of a testing center, proctoring company, or public facility, such as a public library. All proctors must be approved by the faculty. This policy works well in minimizing academic dishonesty and safeguarding academic integrity. However, this is a time-intensive process because each student may have a different proctor who needs to be approved.

Recently, however, there are commercial proctoring services that have become available, and we have had success in having students proctored by them. Some of the companies use human proctors, while others use artificial intelligent agents and computer vision. For example, if a camera identifies that an exam taker constantly has eye or gaze movements in a specific direction, a flag is raised, and a human proctor intervention is triggered.

The scale of online offerings around the world, precipitated by the pandemic, will continue to drive the development of improved technology-assisted proctoring systems. The benefit of working with commercial proctoring services is that they simplify the process while protecting academic integrity. Again, the pandemic has contributed to the mushrooming of products from companies that operate in this domain.

Student engagement and student/faculty interactions have always been perceived as challenging for online courses. Fortunately, with the advances in video conferencing technology, this issue is much less of a concern now. There are many products on the market that allow for synchronous meetings and the conducting of office hours where instructors answer questions and demonstrate how to solve homework problems in real time and have sessions recorded for students who cannot attend. When the meetings with several students are in real time, some of these products offer the use of break-out groups to allow for collective discussion among students. Later, the students are brought back to the meeting for further discussion. These activities promote student engagement.

While the online offering of courses possesses challenges as outlined, there are some unexpected benefits and triumphs. By and large, the online offering mode provides students with a more studious experience in which they can have more opportunities to reflect on the lectures and conduct thoughtful engagement. For example, through the use of online discussions and various communication tools, students have more time to have meaningful exchanges with peers, teaching assistants, and instructors. For students who are more reserved, the online environment is more conducive to discussion as it gives them time to construct responses and questions. While student involvement in online courses is often considered challenging, we found that, if the right tools are used, the online platform can be quite effective for learning.

At the same time, the online platform facilitates discussions and collaboration among cross-disciplinary communities, free from the confinements of time and space. Some of our most successful senior design projects have been carried out by teams of traditional undergraduate students in the face-to-face program and nontraditional students in the BSEE program. The youth and energy of the former and the maturity and experience of the latter have created dynamic bonds on the teams, which has led to very positive project outcomes.

The online program has several advantages. Two of them are rather important. One is the reduced cost of obtaining a bachelor's degree. Clearly, the costs of traveling and living on campus are eliminated as are various fees that onsite students have to pay. The other advantage is that the online delivery of the program allows for reaching out to a very large pool of prospective students around the world.

Our vision

Interestingly, the changes we experienced in academia in 2020 due to the pandemic have helped our program and solidified our intentions to strengthen it further. First, we find ourselves in a situation where we can also learn from many others who have been pushed to offer online courses, in particular courses with strong laboratory components. Second, the market is now much richer with products that allow for online education. This not only includes products used for conducting laboratory exercises but also products for video conferencing and proctoring. With these developments, we find ourselves in a position where we can further enrich our program with laboratory offerings and strengthen every single course by using better technology.

We are mindful of the fact that on-campus classes offer students something that online students do not get, such as the social and cultural experiences of living with their peers. However, we keep in mind that our program is primarily for nontraditional students whose sets of challenges are different from those of the on-campus students, and the BSEE program is optimized for them. Thus, we continue pursuing our aspirations of continued enhancement of the program. We combine these efforts with attempts to attract an increased number of students who will find our program an excellent opportunity for achieving their professional goals.

We envision that in 10 years, our program will become a model program for serving nontraditional students as well as special groups of students who are unable to attend on-campus courses. Following the ABET spirit for continual

By and large, the online offering mode provides students with a more studious experience in which they can have more opportunities to reflect on the lectures and conduct thoughtful engagement.

improvements, the program will expand its technical offerings to include more hands-on components and technical electives and will most likely offer specializations that are driven by the needs of industry, such as machine learning and data science. Further, our department is currently in the process of establishing an online master of science program in engineering artificial intelligence. Once approved by the New York State Education Department, we plan to create an accelerated program for our BSEE students, where, as the name suggests, all of the requirements for completing the degree could be fulfilled more quickly, e.g., by allowing some of the graduate courses in the master's program to be used as technical electives of the BSEE program. We also envision that our program will inspire the creation of similar programs for serving those who are unable to attain a bachelor's degree in the traditional way.

In summary, the journey of the BSEE program has been equally enriching for students and faculty. Motivated by the dedication and preservation of nontraditional students, the instructors have put genuine efforts into delivering online courses of high value. In the process, they have also become learners in the journey of continuous improvement of the BSEE—thus enabling the "learning to learn" paradigm, where the teachers also become the learners.

Authors

Wendy Tang (wendy.tang@stonybrook.edu) received her Ph.D. degree in electrical engineering from the University of Rochester. She has served as a faculty member in the Electrical and Computer Engineering Department at Stony Brook University for more than 20 years. She is the program director for the Bachelor of Science in Electrical Engineering Online program and the associate chair of the Department of Electrical and Computer Engineering at Stony Brook University, Stony Brook, New York, 11794, USA. She has been active in science, technology, engineering, and mathematics education and has been the director/associate director for several National Science Foundation-funded projects with an education focus. Her research interests are in wireless sensor networks, graph theory applications, and communication networks. She is a Member of IEEE.

Petar M. Djuric (petar.djuric@stonybrook.edu) received his B.S. and M.S. degrees in electrical engineering from the University of Belgrade and his Ph.D. degree in electrical engineering from the University of Rhode Island. He is a State University of New York Distinguished Professor and currently a chair of the Department of Electrical and Computer Engineering at Stony Brook University, Stony Brook, New York, 11794, USA. His research has been in the area of signal and information processing. He was a recipient of the European Association for Signal Processing (EURASIP) Technical Achievement Award in 2012. He was the first editor-in-chief of *IEEE Transactions on Signal and Information Processing Over Networks*, and he is a Fellow of IEEE and EURASIP.

References

[1] J. Bourne, D. Harris, and F. Mayadas, "Online engineering education: Learning anywhere, anytime," *J. Eng. Educ.*, vol. 94, no. 1, pp. 131–146, 2005. doi: 10.1002/j.2168-9830.2005.tb00834.x.

[2] "Reimagining the role of technology in education." U.S. Department of Education, Washington, D.C. 2017. https://tech.ed.gov/files/2017/01/NETP17.pdf

[3] F. Mayadas, "Asynchronous learning networks: A Sloan foundation perspective," *J. Asynchro. Learn. Netw.*, vol. 1, no. 1, pp. 1–16, 1997.

[4] T. G. Cleaver and L. M. Elbasyouni, "Student online assessment behaviors," *IEEE Trans. Educ.*, vol. 48, no. 3, pp. 400–401, 2005. doi: 10.1109/TE.2004.842886.

[5] "Evaluation of evidence-based practices in online learning: A meta-analysis and review of online learning studies." U.S. Department of Education, Washington, D.C., 2010. https://www2.ed.gov/rschstat/eval/tech/evidence-based-practices/finalreport.pdf

[6] L. Palma, R. F. Morrison, P. N. Enjeti, and J. W. Howze, "Use of web-based materials to teach electric circuit theory," *IEEE Trans. Educ.*, vol. 48, no. 4, pp. 729–734, 2005. doi: 10.1109/TE.2005.850711.

[7] R. T. Abler and I. G. Wells, "Distributed engineering education: Evolution of the telecollaboration stations for individualized distance learning," *IEEE Trans. Educ.*, vol. 48, no. 3, pp. 490–496, 2005. doi: 10.1109/TE.2005.849757.

[8] W. G. Hurley and C. K. Lee, "Development, implementation, and assessment of a web-based power electronics laboratory," *IEEE Trans. Educ.*, vol. 48, no. 4, pp. 567–573, 2005. doi: 10.1109/TE.2005.856147.

[9] J. Gleason, "Teaching mathematics online: A virtual classroom," *J. Online Math. Its Appl.*, May 2006. [Online]. Available: https://www.maa.org/press/periodicals/loci/joma/teaching-mathematics-online-a-virtual-classroom-abstract

[10] J. Engelbrecht and A. Harding, "Teaching undergraduate mathematics on the internet," *Educ. Stud. Math.*, vol. 58, no. 2, pp. 253–276, 2005. doi: 10.1007/s10649-005-6457-2.

[11] R. S. Adams, J. Turns, and C. J. Atman, "Educating effective engineering designers: The role of reflective practice," *Design Stud.*, vol. 24, no. 3, pp. 275–294, 2003. doi: 10.1016/S0142-694X(02)00056-X.

[12] S. Trines. "Educating the masses: The rise of online education in Sub-Saharan Africa and South Asia." WENR. 2018. https://wenr.wes.org/2018/08/educating-the-masses-the-rise-of-online-education

[13] "Curriculum." Stony Brook. https://www.stonybrook.edu/commcms/eeonline/academics/curriculum.php

[14] W. Tang, C. Westgate, P.-L. Liu, and M. Gouzman, "Online-BSEE (online Bachelor of Science in electrical engineering): An asynchronous online electrical engineering degree program with laboratory," *J. Educ. Technol. Syst.*, vol. 42, no. 4, pp. 355–368, 2014. doi: 10.2190/ET.42.4.c.

[15] G. Siemens, "Massive open online courses: Innovation in education," in *Open Educational Resources: Innovation, Research And Practice*, R. McGreal, W. Kinuthia, and S. Marshall, Eds., Vancouver, Canada: Commonwealth of Learning, 2013, pp. 5–16.

[16] R. Rivard, "Three's company: Stanford University, birthplace of two MOOC companies, decides to work with a nonprofit started by MIT and Harvard," Inside Higher Ed, April 3, 2013. [Online]. Available: https://www.insidehighered.com/news/2013/04/03/stanford-teams-edx

[17] L. Pappano. "The year of the MOOC." NY Times. 2012. https://www.nytimes.com/2012/11/04/education/edlife/massive-open-online-courses-are-multiplying-at-a-rapid-pace.html

[18] D. Shah. "By the numbers: MOOCs in 2019." Class Central. 2019. https://www.classcentral.com/report/mooc-stats-2019/#:~:text=Now%20in%20its%20eighth%20year,online%20degrees%2C%20and%20170%20microcredentials.&text=For%20Class%20Central'%20complete%20analysis%20kee%20reading

[19] H. Ledwon. "55 legit Master's degrees you can now earn completely online." Class Central. 2020. https://www.classcentral.com/report/mooc-based-masters-degree/

SP

John E. Ball, M. Jean Mohammadi-Aragh,
and Christopher Goodin

Engaging Students in an Automotive Autonomy Sensor Processing Class

Incorporating active learning and high-fidelity, physics-based autonomy simulation into class projects

©SHUTTERSTOCK.COM/LUCKYSTEP

Modern signal processing (SP) classes should provide a balance between theory and application as well as use active learning exercises to engage students and facilitate learning. A new sensor processing course, Sensor Processing for Autonomous Vehicles (SPAV), was designed with two specific objectives: 1) to successfully engage students using active and collaborative learning and 2) integrate a state-of-the-art, physics-based autonomy simulator into the class.

The course was delivered to local and asynchronous distance students in spring 2020 at Mississippi State University (MSU). The MSU Autonomous Vehicle Simulator (MAVS) was used in the class. We also utilized three miniprojects to bring together theory and practice. We evaluated the course through student feedback. Results indicated that students viewed active exercises and the simulator as beneficial and useful, with multiple students describing those aspects as their favorite part of the course. Nearly all students (39 of 40) reported that they were engaged in the course.

Background

"Signal processing" is defined by the IEEE Signal Processing Society as ". . . the enabling technology for the generation, transformation, and interpretation of information" [1]. Herein, we also consider SP in a broad context, not just the traditional sampled discrete-time series data processing. For instance, deep learning (DL) image processing as well as radar and lidar object detection all come under the general SP umbrella.

Traditional digital SP (DSP) classes often are very math intensive and focus on "traditional" approaches, such as Fourier-based processing, filter structures and design, and so on. There is still a strong need to teach the fundamentals of DSP given its ubiquitous nature, yet there is also the need for classes to expose students to modern data-driven methods; current research trends; industry challenges; and opportunities in diverse spaces, such as image processing, time series processing, and lidar point cloud processing, among others. DSP classes can be rigorous, yet they can also have application- and system-level content.

Digital Object Identifier 10.1109/MSP.2021.3054327
Date of current version: 28 April 2021

Studies show that, in general, students learn better when they actively participate in meaningful learning experiences [2]. Instructors should ask many questions [2, pp. 31, 85–86] to engage students and allow them to obtain a deeper understanding. Active learning has various forms that are useful for college classrooms, including, for example, informal group discussion, think–pair–share, and minute papers. An overview of active learning can be found in [3] and [4], with a catalog of techniques in [5] and a discussion of opportunities and challenges for active learning in computing courses in [6].

Many research efforts have shown the effectiveness of active learning [3], [7], [8] and student-centric learning [9] in engineering classes. In addition to encouraging active participation, instructors should focus on creating meaningful learning situations for students. One way to make learning meaningful is to include real-world applications in the classroom. Examples and activities based on real-world applications support learning by being memorable, sparking interest, helping students connect new information to their prior understanding, and correcting misconceptions [10].

Class development

In recent years, there has been significant research and development into autonomous vehicles (AVs). The Center for Advanced Vehicular Systems (CAVS), a research center at MSU, performs AV research and development with cameras, radars, and lidars. Camera systems are ubiquitous in automotive autonomy due to high-resolution imagery, high data rates, and low cost. However, they struggle in rain, fog, and low-light situations. Radar and lidar are active technologies and, thus, allow more robust operation in fog, rain, and snow compared to red, green, blue (RGB) cameras.

Radar and lidar are fundamentally different in the data they capture and how they are processed. After working with many students over several years at CAVS in AV processing, it was apparent students had fundamental knowledge gaps in areas such as camera calibration and radar/lidar processing and that classes that addressed these gaps were needed. MSU offers multiple classes on machine learning (ML), neural networks, visualization, (traditional) DSP, image processing, and radar. However, prior to the spring 2020 semester, there was not a comprehensive introductory class on AV sensor processing.

This article outlines a special-topics class, SPAV, that was developed and delivered to 37 on-campus and 11 distance students at MSU during the spring 2020 semester. The SPAV class was designed as a 3-h, split-level (senior/master's degree-level) course focusing on sensor processing methods for cameras, radars, and lidars. To facilitate asynchronous distance students, the class did not require any hardware or equipment other than a Windows laptop (already required for all students by the MSU College of Engineering).

The Spav class was designed as a 3-hour, split-level (senior/master's degree-level) course focusing on sensor processing methods for cameras, radars, and lidars.

The class was listed as an electrical and computer engineering class, but students from any major could enroll. The prerequisite for the class is passing a junior-level signals and systems class or instructor consent. The overarching course goal was to provide a highly interactive course with both a breadth and depth of coverage in automotive autonomy topics, including terminology, the current state of the industry, state-of-the-art processing methods, strengths and weaknesses of each sensor modality, general autonomy frameworks, and control strategies and methods.

Class objectives

Objective 1: Successfully engage students using active and collaborative learning

In the course, students were given a high-level overview of how each sensor modality operates, participated in detailed discussions of the strengths and weaknesses of each sensor modality, and discussed state-of-the-art methods for sensor performance evaluation. Examples of course concepts discussed include the following:

- A camera requires demosaicing to get color imagery, calibration to correct for imperfections, and coregistration to align its images with other sensors.
- Radars operate in all-weather conditions and are good at ranging and velocity estimation, but current-generation radars do not provide high-resolution imagery.
- Lidars provide dense and rich 3D data and are good at object localization, but they are color-blind since they operate at one wavelength.

To get students involved and enhance learning, both active and collaborative learning exercises were heavily utilized in the class. Each class session incorporated two or three active or collaborative exercises, although there was one session with a data collection exercise that lasted about 45 min. These exercises engaged students, illustrated class material, helped students learn fundamental concepts, and allowed the instructor and other students to monitor the observations and conclusions of each student or student group. Since the class also had asynchronous distance students, discussion board questions were utilized to facilitate student interactions among local and distance students as well as allow the distance students to participate in active learning exercises.

Objective 2: Utilize a state-of-the-art, physics-based autonomy simulator in the class

It is widely known that complex systems like AVs require not only real-world driving tests but also simulations to provide effective testing and cover rare edge cases [11]. Baraniuk and Padgett state that using interactive simulations provides an environment where students can explore and learn [12]. Students often relate well to visual-based simulations, especially when they can change parameters and see how the results change.

Initially, three potential simulators were examined for inclusion in the SPAV class: MAVS, a noncommercial, open source software library for simulating autonomous ground vehicles; Autonomous Navigation Virtual Environment (ANVEL); and Car Learning to Act (CARLA), an open source simulator for urban autonomy research [13]. Quantum Signal, ANVEL's developer, was purchased by Ford, and the simulator is now not available for general use. Both MAVS and CARLA provide a Python application programming interface (API) and control over weather, sensors, simulation parameters, agents (vehicles and pedestrians), and so on. However, CARLA requires Linux, and most students in the class do not have Linux machines. For these reasons, MAVS was chosen to be the AV simulator and is discussed in detail in the "MAVS" section.

The primary intent of MAVS is to serve as a software library for simulating the terrain, environment, sensors, and vehicle in autonomous navigation.

Class organization and content

The class met twice a week for 75 min per meeting. The instructor advised the class to spend 1–2 h outside of class for every hour spent in the lecture. The class was organized into seven modules, each having a specific focus, and are summarized in Table 1. To assess students' progress, the class had one homework assignment for each module, three miniprojects, two exams, and a final exam.

In module 1, the students were introduced to automotive autonomy, discussed the Society of Automotive Engineering autonomy levels [14], and examined several car models to assess their autonomy levels. Module 2 gave time for the students to install the required software tools: Anaconda, Tensorflow CPU, numpy, and matplotlib for Python; MAVS; and You Only Look Once (YOLO) [15]. Several after-hours sessions were also provided to help students install the required software tools.

The fundamentals of DL were covered in module 3, including deep convolutional neural network (CNN) building blocks. This module also addressed estimating the number of parameters in each layer, which is important for embedded applications. Students also ran a CNN version of MNIST using Tensorflow and learned how to write Tensorflow code.

Module 4 covered decision, planning, and control. In this module, a proportional-integral-derivative (PID) controller was introduced, and the students examined how changing the PID controller parameters affected the response. MPC was discussed in the context of path planning. The students had an exercise where they used MPC to enable a vehicle to avoid obstacles and successfully reach the destination. Finally, reinforcement learning was introduced, and they played a simple game, stepping through the reinforcement learning system as it learned to play.

Module 5 focused on camera processing and started with a discussion of the human eye and how cameras operate in a similar manner to human rods and cones. This module then covered the basics of image demosaicing, the Bayer filter, the pinhole camera, coordinate transformation, camera calibration, stereo processing, and structure from motion. Several state-of-the-art methods were examined. Thermal and infrared (IR) cameras and how they might be used in autonomy were explained. During this module, a FLIR Systems Automotive Development Kit (ADK) (https://www.flir.com/products/adk/) long-wave IR (LWIR) camera was demonstrated to the class. The final lecture in this module was devoted to thermal imaging, and the class reviewed results from studies dealing with thermal cameras [16] to understand issues facing regular RGB cameras and how different types of thermal cameras can help in poor weather conditions.

Module 6 focused on radar processing and included topics on radar terminology, waveforms, and the radar range equation. Next, frequency-modulated continuous-wave (FMCW) radar signaling and processing were covered, including range estimation, range resolution, and maximum range calculations. The class then discussed the specifications of several automotive radars. Finally, Kalman filtering was covered and discussed in the context of adaptive cruise control.

Module 7 covered lidar processing. As most students had no previous experience with lidars, lidar architectures were explained, as were lidar terminology and design parameters (the number of beams, frame rate, maximum object range, and so on). Laser emitters, laser beam divergence, and laser detectors were discussed. Time-of-flight calculations, the lidar range equation, and atmospheric effects on

Table 1. A summary of the class modules and learning objectives.

Modules	Learning Objectives
1: Autonomy	Discuss and explain autonomy levels and basic autonomy modes
2: Tool Install	Install MAVS, Anaconda, and Python tools and utilize them in class
3: DL	Utilize DL to run advanced driver assistance systems (ADAS) processing algorithms Discuss and evaluate state-of-the-art processing methods for radar, lidar, and cameras
4: Control	Utilize proportional-integral-derivative and model predictive control
5–7: Camera, Lidar, and Radar	Explain the capabilities and limitations of radar, lidar, and camera systems Process and analyze results from real-world and simulated autonomy data sets Discuss and evaluate state-of-the-art processing methods for radar, lidar, and cameras Understand and implement basic processing steps for radar, lidar, and camera data Understand the strengths and weaknesses of radar, lidar, and camera ADAS processing

lidar were discussed. State-of-the-art methods in object detection, free-space mapping, and road detection were discussed. A Velodyne HDL-32 and VeloView (https://www.paraview.org/veloview/) were demonstrated to the class. Several scenes were captured using VeloView and, to illustrate that visualizing objects in lidar data is difficult, the students were asked to guess what objects they were seeing.

MAVS

MAVS is an interactive, real-time, physics-based simulator for autonomous ground vehicles [17]. MAVS uses physics-based ray tracing [18] to accurately simulate sensors like lidar and cameras in addition to realistic simulations of GPS sensors and microelectromechanical sensors, such as inertial measurement units and gyroscopes. Vehicle dynamics are simulated in MAVS using ReactPhysics3D [19]; MAVS can also be interfaced with other vehicle dynamics software, such as Chrono [20].

MAVS is free and open source for noncommercial use. The core MAVS libraries are written in C++, and the code can be integrated via the C++ API or Python interface. MAVS is available on GitLab. (It can be downloaded from https://gitlab.com/cgoodin/msu-autonomous-vehicle-simulator.) Additionally, precompiled binaries for Windows 10 and Ubuntu 16.04 (MAVS precompiled binaries can be downloaded at http://www.cavs.msstate.edu/capabilities/mavs.php) and extensive online documentation are also available. (MAVS documentation is available at https://mavs-documentation.readthedocs.io/en/latest/.)

The primary intent of MAVS is to serve as a software library for simulating the terrain, environment, sensors, and vehicle in autonomous navigation. MAVS is structured to either be integrated into other applications or have other software components run in a cosimulation approach. MAVS features four basic simulation modules: vehicles, sensors, environments, and scenes.

The vehicle module provides a simulation of the vehicle motion and dynamics. The scene module defines the geometry, color, and texture of objects within the scene as well as methods for querying scene geometry using ray tracing. MAVS uses several tire and terrain interaction models to simulate driving on a variety of pavement and soil conditions and can simulate a variety of weather and environmental effects and their influences on sensor performance. The impact of rain on lidar in MAVS has been shown to match real measurements [21]. Lighting conditions based on time of day (including night) and atmospheric haziness can also be simulated with MAVS.

MAVS is being used by students, faculty, and staff at MSU to perform research in many areas of off-road autonomous operation including navigation in rough terrain, vegetation and terrain classification, negative obstacle detection, and stop sign detection. The class provided valuable distribution experience and feedback to the MAVS development team in preparation for the public release of the code (https://www.cavs.msstate.edu/capabilities/mavs.php).

For students and researchers studying ML, MAVS can automatically generate semantically labeled data for training and testing ML algorithms. The automated labeling process has been used for testing neural network-based ML algorithms for both camera [22] and lidar [23] data. Some labeled camera outputs are shown in Figure 1.

In addition to using MAVS data, students were also given databases of road scenery collected locally by the instructor and a student containing dirt and paved roads, various signage, and so on. These scenes covered highway, more country-like settings, and some urban (downtown) areas.

For students and researchers studying ML, MAVS can automatically generate semantically labeled data for training and testing ML algorithms.

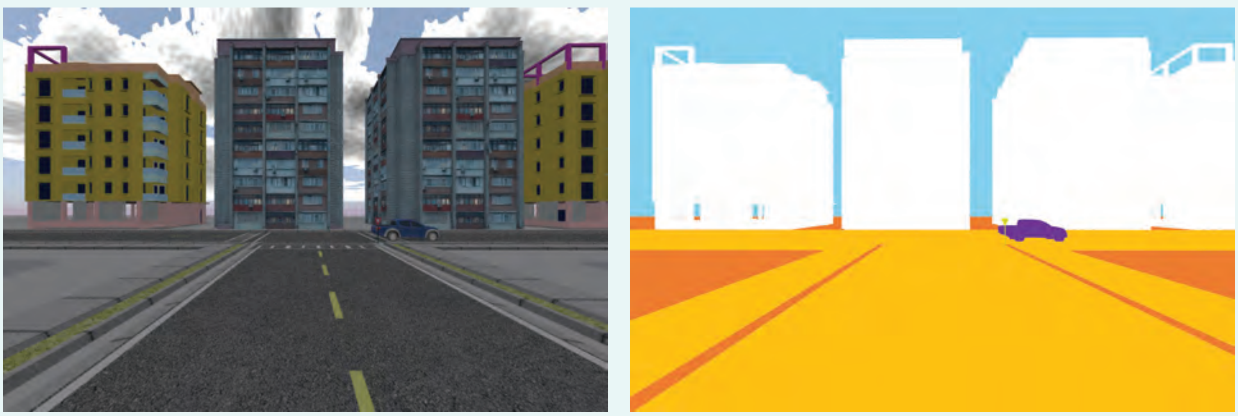

(a) (b)

FIGURE 1. An example of MAVS automatic semantic labeling: the (a) raw and (b) labeled image. The white label is for buildings, purple indicates vehicles, blue represents sky, yellow shows road, and orange is for ground. A yield sign can be seen labeled in light green near the truck. In particular, the extension to noisy microwave networks is discussed in detail with respect to the interface with optimization algorithms, a topic that should attract a wide readership.

Active and collaborative learning exercises

A total of 71 active and collaborative learning exercises were used in the class. Exercises included brainstorming activities, where the goal was to think of as many responses as possible; think–pair–share activities, where each student thought about a problem or question, discussed with the other students in the group, and finally came to a consensus; and various types of group exercises. Exercise lengths ranged from several minutes up to about 50 min. (The class runs for 75 min.) Most of the exercises were shorter and designed to reinforce concepts.

An example of a shorter exercise was listing challenges to implementing level 5 autonomy. Example student responses are edge cases, price, ethics, handling construction, working with drivers in level 0 vehicles, malfunctions, handling aggressive drivers, and so on. This exercise took about 10 min.

There were a variety of collaborative exercises involving examining and running DL or sensor processing code.

There were also more in-depth exercises in the class. Some examples include the following:

- having student groups take processing steps, e.g., mapping, localization, traffic prediction, and so on; explain where these tasks fit into the Eliot artificial intelligence (AI) automotive framework (a block diagram for an autonomous system); and give their rationale [24] (20 min)
- a detailed analysis of a radar Blake sheet for an FMCW radar: an Excel spreadsheet was handed out, and students investigated how certain parameters affected the radar performance) (20 min)
- a Kalman tracker simulation, where the students examined the effects of changing two parameters in the Kalman filter in a filter simulation (25 min).

In all of these cases, team results were posted to discussion boards.

Inevitably, there are gaps between theory and practice, and many algorithms or methods work well with small or limited data but might have issues in the real world. Several exercises were geared toward exploring these areas. A discussion board exercise asked students to review a paper and discuss potential difficulties encountered with using a lidar in the rain and, in a second exercise, with pedestrian detection in fog (with various types of thermal, IR, and color cameras). A different assignment asked groups of students what difficulties there could be with road detection algorithms, especially considering the many dirt and gravel roads in rural Mississippi, snow-covered scenes in northern states, and flooded areas, to name a few. Another task asked students to consider what happens in the case of a free-space mapper and path planner where there is no free space in front of the vehicle (e.g., following someone or parked in a parking lot).

There were a variety of collaborative exercises involving examining and running DL or sensor processing code. These involved groups of students and were performed in class for the local students. To facilitate asynchronous distance student involvement, the collaborative exercises culminated with the groups posting to discussion boards, where the distance students would also review and post comments in the following few days.

For example, in module 2, students examined code for a CNN to classify the MNIST digit data set. They trained the CNN and ran inference on the testing images. This exercise introduced them to DL and allowed the instructor to explain the basic MNIST CNN. A later miniproject allowed students to investigate using YOLO 9000 [15] for sign detection in simulated and real imagery.

In another instance, student groups ran two Python QT5 GUIs that demonstrated radar SP. The first GUI let them discover that, in FMCW processing, the distance of a reflecting object (we used a point target) after FMCW demodulation results in an intermediate frequency (IF) that is proportional to the object's distance. Instead of first giving them the equation that relates the IF to the object distance, the students ran simulations and hypothesized that, as the distance increases, the IF increases also. They had a visual understanding, and then we confirmed that their hypothesis is correct and that there is a linear relationship between the IF and object distance.

The second GUI gave insight into FMCW radar processing, and it allowed students to visualize automotive radar object detection. They could change the radar's FMCW parameters as well as the object's radar cross section, range, and velocity. This GUI is shown in Figure 2. The top plot on the right shows the range fast Fourier transform (FFT) results, and the bottom right plot shows the range–velocity results after velocity FFT processing. The class discussed the relationship of the IF to the object distance from the radar.

Other collaborative exercises focused on system-level information. For example, in the radar module, students used a spreadsheet and modified radar parameters for a short-range radar. When specifications were met, cells turned from red to green. Also in the radar module, students ran a Kalman filter simulation and tuned the filter parameters to see how they affected the results. In the lidar module, students examined a lidar design that calculated the maximum lidar frame rate given the field of view (FOV), number of pulses, pulse widths, and number of receivers.

In the decision, planning, and control module, students listed challenges to an autonomous system as a vehicle approaches an intersection; they also took a set of software modules defined in [25] and mapped them into the Eliot automotive framework [24]. Students were asked to explain their choices in this exercise.

Miniprojects

The classwork included three miniprojects assigned by the instructor. In each of these, local and distance students worked in teams of up to four undergraduates or four graduates (with no mixed teams). Each miniproject required the software-based assignment to be conducted. Each team submitted a report with an introduction, methodology, results, conclusions, references, and code listings. The miniprojects were designed to teach how to perform and write a small research project. Grading was based on following directions, technical content, proper IEEE citations, grammar, profession writing style, and code comment clarity. Each miniproject was worth 10% of the final grade. The timeframes were five, five, and four weeks for miniprojects 1–3, respectively.

The first miniproject allowed students to run MAVS for the first time and utilize a pretrained YOLO 9000 [15] to allow them to see how well a state-of-the-art detector would work to detect stop and yield signs in high-fidelity simulated driving imagery. Students performed experiments and wrote their results in a final report for each miniproject. Figure 3 shows example MAVS imagery.

The second miniproject was given after students had discussed camera operation and learned about camera calibration as well as camera model intrinsic and extrinsic matrices. Students collected data in class, and student groups performed an offline camera calibration procedure with the full data set and a partial data set. They then examined the calibration results and wrote a report on their findings. They discovered that you need a variety of poses and you must have samples all around the camera FOV to obtain a good calibration.

Figure 4 shows three students collecting camera calibration data in class. Distance students participated in all exercises. In the camera calibration exercise, they were not able to collect data; however, they posted their observations of how well the in-class students performed the calibration data collection, e.g., whether they got images covering a variety of the image space, different orientations of the calibration board, and so on.

Perhaps the most engaging for students was the third miniproject. Most students in the class had no experience with lidar and lidar processing. After learning about laser emitters, laser detectors, scanning lidars, and so on, they used MAVS to simulate a lidar detecting a brick on the road. The simulations examined the following lidars: Velodyne VLP-16, HDL-32E and HDL-64E; Ouster OS1 and OS2; and a Quanergy M8. The simulation estimated the number of lidar points reflected from a brick at given distances from the vehicle. The students studied how the different lidars would behave.

Class assessment: Challenges faced

There were many challenges in preparing and administering the class. Most instructors who have had to prepare a class for the first time will agree that this is a daunting task by itself. The first challenge was the depth versus breadth of the class. We wanted it to not only contain sufficient depth but also breadth as well as to focus a majority of the class on sensor processing methods for the lidar, camera, and radar sensors. To prepare students for state-of-the-art discussions, which mostly involved DL methods, an early module on DL

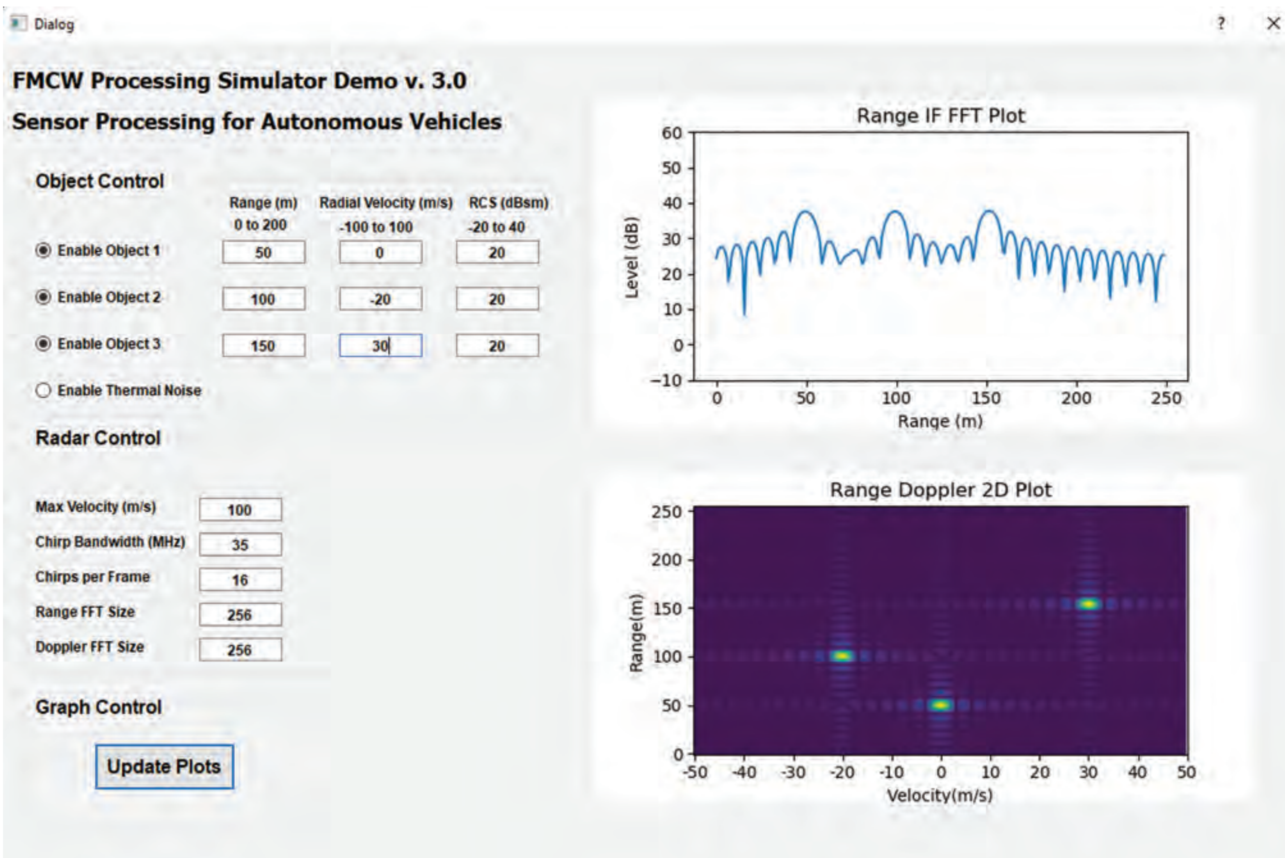

FIGURE 2. The radar FMCW processing GUI. FFT: fast Fourier transform; Max: maximum. RCS: radar cross section; IF: intermediate frequency; FFT: fast Fourier transform.

was created. Though a single class cannot cover all topics, the idea was to briefly explain major topics, such as camera intrinsic and extrinsic matrices, and expose the class to image processing topics, such as camera calibration, stereo processing, and structure from motion.

The second challenge was that there was no single book that covered the material. For a traditional DSP class, there are myriad books available. Traditional DSP is a very mature field, while automotive autonomy is changing rapidly and still in a developmental phase. Three books were selected. The first was *Creating Autonomous Vehicle Systems* [25], which covered autonomy in general; localization; perception; prediction and routing; decision, planning, and control; and reinforcement learning. To cover autonomy complexity, system framework, graceful degradation, ML, ethical issues, and so on, Eliot's book *Introduction to Driverless Self-Driving Cars: The Best of the AI Insider* [24] was chosen. *Computer Vision in Vehicle Technology: Land, Sea, and Air* was chosen to cover computer vision and vision-based autonomy systems [26].

These books did not provide adequate coverage of radar and lidar. The class materials and supplemental journal articles were used to cover these topics. We note that a very good book on autonomous radar, *Radar Signal Processing for Autonomous Driving* [27], was published too late for our course offering, but we will use the book in future classes as it is written by a nonradar expert aimed at other nonradar experts.

Covering state-of-the-art methods meant students had to read journal papers. Most graduate students are accustomed to doing this, but undergraduates are not. Having all of the students select and critique papers in a one-page writing assignment as part of each module homework gave students experience with literature reviews, how to scan a paper to find the key concepts and contributions, and how to effectively write a critique of the paper's pros and cons.

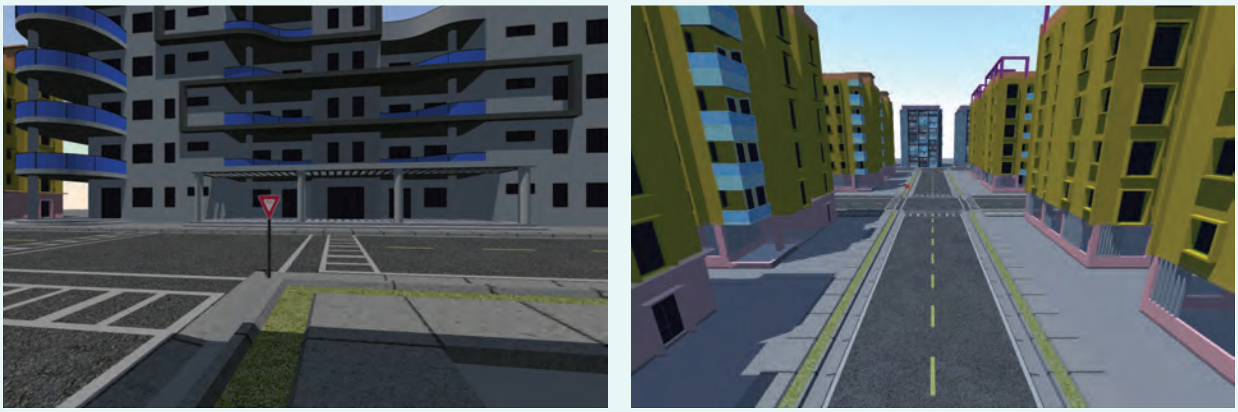

(a) (b)

FIGURE 3. The MAVS-generated urban scenes used for class miniproject 1: (a) a yield sign and (b) a four-way intersection.

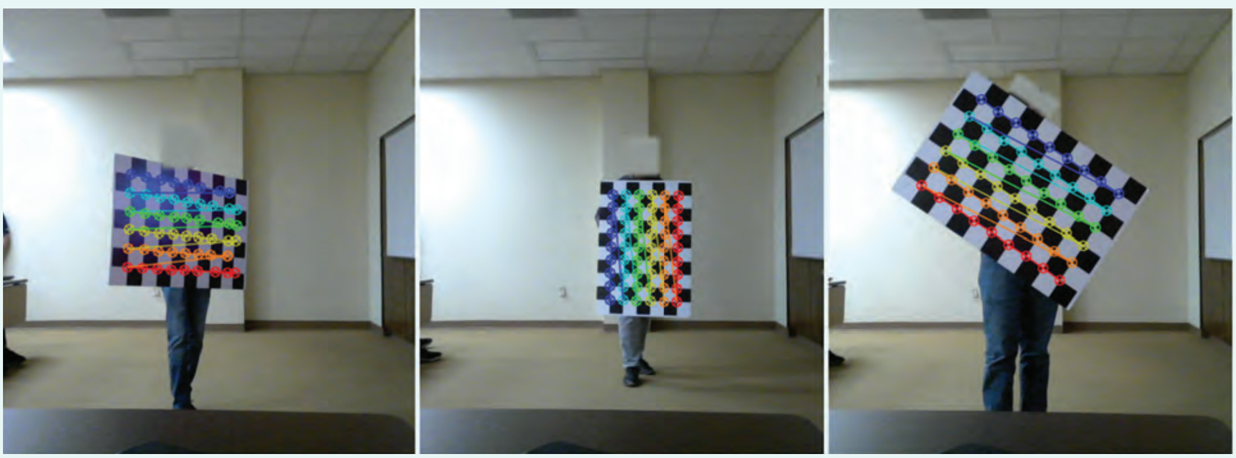

FIGURE 4. The miniproject 2 camera calibration in-class exercise.

A third challenge is the sheer amount of materials required for the class: lectures, papers for the students to read and critique, and Python codes for in-class demos. Both the DL and autonomy fields are changing rapidly, and, since there were about six class days devoted to state-of-the-art methods, these sections will need to be revised each semester as new techniques overtake the older approaches or existing methods are updated and significantly improved.

A fourth challenge was asynchronous distance teaching. Distance students often have a highly varied background, are usually working full time, and, often, have families and other duties. Most distance students work asynchronously, so they lag behind the local students since they usually watch videos at night or on weekends. Distance students also do not have the benefit of working directly with other students, unless there are several distance students who work at the same company. Keeping distance students engaged and having them feel involved is very challenging. We believed that utilizing discussion boards and having mixed groups (distance and local students) on the miniprojects helped to keep them involved.

The final—and very much unexpected—challenge was the COVID-19 pandemic, which moved all MSU post-spring break classes from in person to online. Since the course materials were organized as PowerPoint presentations planned for both in-person and distance offerings, the challenges with the transition to fully online classes were somewhat mitigated. The course instructor (the first author) had never taught an online class. The distance class was taught in a special classroom and recorded for distance students to participate asynchronously. After the COVID-19 transition, a majority of the local students participated synchronously, with the active learning exercises continuing.

Several approaches changed as the class met online:

1) Before spring break, the instructor would annotate materials using the SMART board display in the distance learning classroom. For the online class, the instructor utilized a second camera and wrote on paper. Students could see the writing, and scanned versions were distributed after the class.
2) After spring break, the instructor started using WebEx Polls to poll students.
3) Most students do not prefer to interact remotely with videos on, so the instructor could not see most of the students. Before the online class, the instructor would walk around and talk to students about the exercises and provide feedback.

Student feedback

The course was designed with the two specific objectives: 1) to successfully engage students using active and collaborative learning and 2) integrate a state-of-the-art, physics-based autonomy simulator into the class. We evaluated the course through student feedback, which provided their perceptions of the course. Students provided informal feedback during the entire semester as a regular part of the active learning exercises. During the semester, there were three opportunities for students to give formal feedback to the instructor: an anonymous survey, a bonus question on the final exam, and the standard university-administered course evaluation. In this section, we discuss results from the formal survey and final exam question.

Final exam bonus question feedback

The final exam for the course included an open-ended "bonus" question that prompted, "What did you like the most about this class?" All student feedback via the bonus question was favorable. With regard to objective 1, students appreciated the active learning exercises for forcing engagement with the course topics. As an example, one student commented,

> What I liked most about this class was that you forced the class to be involved. It is easy just to sit and "attend" a lecture, but you made it fun and interactive. I also think that the class exercises were a huge help. I loved that we were able to solve the problems in class instead of only working problems at home and being lost.

Another student highlighted that the active exercises, which included both demonstrations and tinkering, helped solidify course concepts:

> My favorite part of the class was the demonstration of the various sensors and seeing them work in real time. Specifically, the in-class taking of camera calibration images, live demo of the thermal camera, and live lidar mapping of your office. To be able to visualize the output of the sensors is critical to an intuitive understanding of a sensor system. Second to that, I liked the projects that showed us the output of the camera calibration and radar display programs. Playing with parameters and seeing the effects is very satisfying.

During the lectures, students were responsive to the active learning exercises. At the conclusion of the course, numerous students described the active exercises as their favorite part of the class.

With regard to objective 2, student responses reflective positive perceptions of integrating real-world applications into the course. As expected, students highlighted how real-world applications helped them translate the theoretical course concepts to specific engineering contexts. For example, one student said,

> The combination of theory and application is what every engineering course should consist of. This class purely shows your expertise in the field, and you have the ability to hand down parts of that knowledge to us. . . . The books and articles were nice to be able to read and interpret. Getting exposure to Python, Anaconda, and MAVS are all transferable skills to take us to the next level of expertise within the field.

Another student agreed that the real-world applications enhanced the course:

At the conclusion of the course, numerous students described the active exercises as their favorite part of the class.

I enjoyed the projects a lot. I like being able to see what we learn in class and being able to apply it in the real world. I am a very applicable and applied person, so I enjoy the applying side of the class more. Like I said, the projects allowed me to see and use what we learned in class works. I also really enjoyed learning the different ML algorithms and the image processing to track objects—that was super cool. I love how we can run these algorithms, and a computer can be trained to look at live images and pick out specific things to track or tell us is there with very high precision.

One student discussed that the real-world activities help them connect and understand fundamental SPAV course concepts with ideas learned in other undergraduate engineering courses:

My favorite part about this class was learning how to apply everything I've learned in my four years of undergraduate study. This class took everything from Python code, linear algebra, circuits, and signal processing put a major application on it, the vehicle self-driving vehicle. It has given me a lot of appreciation towards where the vehicle self-driving vehicle industry is and where it will go.

Additionally, multiple student responses specifically mentioned the benefits of using the state-of-the-art, physics-based autonomy simulator in class. They appreciated that the MAVS software was currently used to solve automotive autonomy problems—for example,

What I liked most about the class were the simulations with MAVS, demos, the Python executable codes, and the grad project. For miniproject 1, it was neat knowing that real applicable simulations could be executed with MAVS and its data could be valid to further develop approaches in automotive autonomy. The demos, such as the camera calibration, lidar point cloud analysis, and the FLIR thermal camera showcase, were very interesting.

By recalling specific aspects of the MAVS projects, such as changes to stop signs, one student indicated that the projects were memorable and achieved the goal of creating meaningful active learning:

My favorite part of the class was working on the miniprojects, especially with MAVS. The simulation of the ground vehicle was very fascinating to me, and I had a lot of fun interacting with the different variables and changing the vehicle paths and the environment variables. It was very interesting to see how slight changes to variables could greatly affect the image quality of the stop signs.

Numerous students specifically mentioned MAVS when describing their favorite aspects of the course. Several students noted that is was helpful to be able to have "hands-on" experience that applied the concepts they learned in class. It also seemed that the segmentation project using MAVS with YOLO was popular because of the visual nature of the algorithm.

The ideas of using active and collaborative learning, incorporating simulations, discussing state-of-the-art methods, and using miniprojects can be incorporated into many engineering classes.

Survey assessment

In addition to the exam question, an anonymous Qualtrics survey was administered. The questions are listed in Table 2. Of the 48 students enrolled, 40 responded to the survey, for a response rate of 83%. Overall, student responses indicate that the coverage of state-of-the-art methods and DL was beneficial (question 1). With regard to objective 1, 39 of the 40 students who responded to question 4 agreed that they felt engaged, which was a major goal for both the synchronous distance and local students. All responses indicated that students enjoyed hardware demos (question 2), and 38 of 39 students reported that the hardware demos improved their understanding of the course concepts.

Questions 6–8 in Table 2 focused on the usefulness of the active exercises (objective 1) and the simulator (objective 2) for creating a meaningful, engaged learning experience. Thirty-five of 39 students reported that the active and collaborative learning exercises were extremely or very useful for learning. When specifically asked about the MAVS software, 30 of 40 students viewed the software as useful for illustrating concepts and performing experiments (question 7), and 29 of 39 viewed MAVS as useful for learning in the general context. A few students indicated that MAVS was not useful.

In addition to the questions in Table 2, the survey included an optional open-ended question: "Briefly provide any reasoning for your views of the usefulness of visitors, student exercises, or the MAVS simulator for learning." In response to that prompt, one student noted that the active exercises and MAVS were more beneficial once students were required to connect concepts and implement automotive autonomy tasks:

The exercises and MAVS always seemed useful during the class, but became clear just how beneficial they are during the final project.

Another student perceived MAVS as useful because it allowed for students to further explore topics beyond the provided course content:

The exercises using tools and simulations really help drive the points home and to allow for experimentation with learned principles outside of class.

Students did not provide reasoning for their unfavorable ratings, but we believe negative views could be related to the specific challenges some of them encountered when using the software. While some students in the class had experience with Python programming, others did not. Additionally, during the course, a few students suggested software improvements, such as increasing the size of the simulation screen and providing more built-in file-type exports. We note that challenges like these when learning new software are not unique to MAVS.

In the course, the instructor tried to strike a balance in theory and applications. Moreover, the class was designed to fit a need for an introductory class that covered the three major autonomy sensors, their operating principles, and writing DL code to understand sensor operations and limitations. Table 2 shows that, overall, students felt that there was a relatively good balance of theory and applications. Of the 40 respondents to that question, 34 said the balance was about right, three opined that the class was a little too theory oriented, and three rated the class as too application oriented.

When designing the course, we viewed theory as a required element because we anticipated the course material would be new to most of the students. Topics such as how cameras capture data and use Bayer filters to generate RGB imagery, how camera intrinsic and extrinsic parameters are defined and how to estimate them, and how lidars and radars operate were all discussed in the class. Our initial assessment of students' prior knowledge was correct, as shown by the responses in Table 2, where 27 students indicated limited prior knowledge (a little or none). Seven students indicated they had a moderate amount of prior knowledge, with six indicating they had a lot or a great deal of prior knowledge. We believe most students' prior knowledge came from work experience at CAVS or a course on radars.

Student perceptions of the course were positive and encouraging. Through their survey responses, students reported that they were engaged in learning, enjoyed the hardware demos, and viewed the course concepts as beneficial. Students also reported that the active exercises and incorporation of the MAVS simulator in class were very or extremely useful (35 out of 39 and 29 out of 39, respectively). Students' behaviors, including class attendance, participation in activities, and posting regularly to the discussion boards, further indicated that they valued the active exercises.

> **"My favorite part about this class was learning how to apply everything I've learned in my four years of undergraduate study."**

Conclusions

A 3-semester-h class, SPAV, offered at the senior/master's degree level, was developed from scratch for MSU. The class focused on automotive autonomy, DL, and sensor processing for lidar, camera, and radar sensors. The class was designed to expose students to the three primary sensor systems in AVs and give them hands-on experience in sensor processing and state-of-the-art methods.

Since this class was offered as a special topics class, it will undergo another revision and offering and then be submitted to the MSU curriculum committee for adoption as a permanent class. It is the intent of the authors to strongly pursue cross-disciplinary enrollment. Currently, any engineering major can take this course as an elective. Emails advertising the new class offering will be sent to all engineering departments and researchers at CAVS so that interested students can have the opportunity to take the class.

This class was challenging to develop, and the authors do not recommend that a pretenured assistant professor undertake a new-start class that is so demanding. However, the ideas

Table 2. A summary of the survey question responses.

Question	*n*	Strongly Agree	Somewhat Agree	Neither Agree or Disagree	Somewhat Disagree	Strongly Disagree
1) The coverage of DL/modern state-of-the-art methods is very beneficial.	39	**29**	10	—	—	—
2) I enjoy the hardware demos.	40	**37**	3	—	—	—
3) The hardware demos help me understand the sensors.	39	**33**	5	1	—	—
4) I feel engaged as part of this class.	40	**33**	6	1	—	—
5) I feel challenged due to new material that is part of this class.	38	**21**	14	3	—	—
		Extremely Useful	Very Useful	Moderately Useful	Slightly Useful	Not at all Useful
6) The student exercises are beneficial (useful) for learning in the class.	39	**23**	12	3	1	—
7) The MAVS simulator is beneficial (useful) for illustrating concepts and performing experiments.	40	**25**	5	6	2	2
8) The MAVS simulator is beneficial (useful) for learning in the class.	39	**21**	8	7	1	2
		Too Theory Oriented	A Little Too Theory Oriented	About Right	A Little Too Application Oriented	Too Application Oriented
9) How well does the course balance theory and applications?	40	—	3	**34**	1	2
		A Great Deal	A Lot	A Moderate Amount	A Little	None at All
10) How much of the material in this class did you know prior to taking the class?	40	3	3	7	**21**	6

The *n* in column two indicates the total number of respondents. The largest response categories are shown in bold font.

of using active and collaborative learning, incorporating simulations, discussing state-of-the-art methods, and using mini-projects can be incorporated into many engineering classes.

Using MAVS in the class was not only beneficial to students but also valuable to the MAVS developers, as it provided a group of testers with a diverse range of experience and technical backgrounds. Students provided excellent actionable feedback for improving MAVS, pointing out the need to make the installation process easier and provide more examples and training.

The feedback on the use of active exercises and incorporation of the MAVS simulator in the class was overwhelmingly positive. Students provided informal feedback throughout the course as part of the active exercises, which was used to hone the classroom experience in real time to strengthen the learning experience.

Students also provided more formal perceptions of the course through the use of a feedback prompt and a 10-question survey. Multiple students described the active exercises or the MAVS simulator as their favorite part of the course. Student perceptions of the usefulness of the exercises and MAVS were nearly all positive. The results demonstrate that the course achieved the objectives of successfully 1) engaging students using active and collaborative learning and 2) integrating a state-of-the-art, physics-based autonomy simulator to create meaningful active learning in the classroom.

Authors

John E. Ball (jeball@ece.msstate.edu) received his Ph.D. degree in electrical engineering from Mississippi State University (MSU). After spending 16 years in several industry and government positions, he is currently an associate professor and the Robert Guyton Endowed Chair of Teaching Excellence in the Electrical and Computer Engineering Department at MSU, Mississippi State, Mississippi, 39762, USA. His research interests include sensor processing, automotive autonomy, deep learning, and effective teaching in engineering. He is a Senior Member of IEEE.

M. Jean Mohammadi-Aragh (jean@ece.msstate.edu) received her Ph.D. degree in engineering education from Virginia Tech. She is an assistant professor in the Electrical and Computer Engineering Department at Mississippi State University, Mississippi State, Mississippi, 39762, USA. Her research interests include learning analytics, pedagogical uses of digital systems, undergraduate engineering degree pathways, and diversity in engineering. She is a Member of IEEE.

Christopher Goodin (cgoodin@cavs.msstate.edu) received his Ph.D. degree in physics from Vanderbilt University. From 2008 to 2017, he worked with the U.S. Army Engineer Research and Development Center in Vicksburg, Mississippi, developing physics-based simulations of ground vehicles, sensors, and robotics. Since 2017, he has been a research faculty member with the Center for Advanced Vehicular Systems, Mississippi State University, Mississippi State, Mississippi, 39762, USA, specializing in the development of autonomous ground vehicle simulations.

References

[1] IEEE Signal Processing Society. 2020. https://signalprocessingsociety.org/ (accessed: Feb. 25, 2020).

[2] K. Bain, *What the Best College Teachers Do.* Cambridge, MA: Harvard Univ. Press, 2011.

[3] P. Baepler, J. D. Walker, D. C. Brooks, K. Saichaie, and C. I. Petersen, *A Guide to Teaching in the Active Learning Classroom: History, Research, and Practice.* Sterling, VA: Stylus, 2016.

[4] R. M. Felder and R. Brent, "Active learning: An introduction," *ASQ Higher Educ. Brief*, vol. 2, no. 4, pp. 1–5, 2009.

[5] J. L. Faust and D. R. Paulson, "Active learning in the college classroom," *J. Excell. College Teach.*, vol. 9, no. 2, pp. 3–24, 1998.

[6] K. E. Holbert and G. G. Karady, "Strategies, challenges and prospects for active learning in the computer-based classroom," *IEEE Trans. Educ.*, vol. 52, no. 1, pp. 31–38, 2009. doi: 10.1109/TE.2008.917188.

[7] M. Prince, "Does active learning work? A review of the research," *J. Eng. Educ.*, vol. 93, no. 3, pp. 223–231, 2004. doi: 10.1002/j.2168-9830.2004.tb00809.x

[8] R. Anderson, R. Anderson, K. M. Davis, N. Linnell, C. Prince, and V. Razmov, "Supporting active learning and example based instruction with classroom technology," in *Proc. SIGCSE 2007, 38th SIGCSE Techn. Symp. Comput. Sci. Educ.*, pp. 69–73. doi: 10.1145/1227310.1227338.

[9] E. Ann, J. Anthony, M. T. Kathleen, D. Janet, D. Cheryl, J. D. Filer, C. D. Wiedmaier, and C. Moore, "Students' perceptions of characteristics of effective College Teachers University of Arkansas at Fayetteville," University of Arkansas, Tech. Rep. 5, 2003.

[10] J. D. Bransford, A. L. Brown, and R. R. Cocking, *How People Learn: Brain, Mind, Experience, and School*. Washington, D.C.: National Academy Press, 2000.

[11] T. Costlow. "Complexity of autonomous-systems simulation, validation soars to the clouds." May 2019. https://www.sae.org/news/2019/06/av-simulation-testing-trends (accessed: May 17, 2020).

[12] R. G. Baraniuk and W. T. Padgett, "Trends in signal processing education [in the spotlight]," *IEEE Signal Process. Mag.*, vol. 29, no. 1, pp. 180–182, 2011. doi: 10.1109/MSP.2011.943132.

[13] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proc. 1st Annu. Conf. Robot Learn.*, 2017, pp. 1–16.

[14] *Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems*, SAE International, SAE Standard J3016_201401, 2014. [Online]. Available: https://www.sae.org/standards/content/j3016_201401/

[15] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," 2016, arXiv:1612.08242.

[16] N. Pinchon, O. Cassignol, A. Nicolas, F. Bernardin, P. Leduc, J.-P. Tarel, R. Brémond, E. Bercier et al., "All-weather vision for automotive safety: Which spectral band?" in *International Forum on Advanced Microsystems for Automotive Applications*. Berlin: Springer-Verlag, 2018, pp. 3–15.

[17] C. Goodin, M. Doude, C. R. Hudson, and D. W. Carruth, "Enabling off-road autonomous navigation-simulation of lidar in dense vegetation," *Electronics*, vol. 7, no. 9, p. 154, 2018. doi: 10.3390/electronics7090154.

[18] I. Wald, S. Woop, C. Benthin, G. S. Johnson, and M. Ernst, "Embree: A kernel framework for efficient CPU ray tracing," *ACM Trans. Graph. (TOG)*, vol. 33, no. 4, pp. 1–8, 2014. doi: 10.1145/2601097.2601199.

[19] ReactPhysics3d. https://www.reactphysics3d.com/

[20] R. Serban, M. Taylor, D. Negrut, and A. Tasora, "Chrono:: Vehicle: Template-based ground vehicle modelling and simulation," *Int. J. Veh. Perf.*, vol. 5, no. 1, pp. 18–39, 2019. doi: 10.1504/IJVP.2019.10018132.

[21] C. Goodin, D. Carruth, M. Doude, and C. Hudson, "Predicting the influence of rain on LIDAR in ADAS," *Electronics*, vol. 8, no. 1, p. 89, 2019. doi: 10.3390/electronics8010089.

[22] C. Goodin, S. Sharma, M. Doude, D. Carruth, L. Dabbiru, and C. Hudson, "Training of neural networks with automated labeling of simulated sensor data," SAE Techn. Paper 2019-01-0120, 2019.

[23] L. Dabbiru, C. Goodin, N. Scherrer, and D. Carruth, "Lidar data segmentation in off-road environment using convolutional neural networks (CNN)," SAE Techn. Paper 2020-01-0696, 2020.

[24] L. Eliot, *Introduction to Driverless Self-Driving Cars: The Best of the AI Insider.* LBE Press, 2018.

[25] S. Liu, L. Li, J. Tang, S. Wu, and J.-L. Gaudiot, *Creating Autonomous Vehicle Systems.* San Rafael, CA: Morgan & Claypool, 2017.

[26] A. M. López, A. Imiya, T. Pajdla, and J. M. Álvarez, *Computer Vision in Vehicle Technology: Land, Sea, and Air.* Hoboken, NJ: Wiley, 2017.

[27] J. Gamba, *Radar Signal Processing for Autonomous Driving.* Berlin: Springer-Verlag, 2020.

SP

Elias Aboutanios, Vidhyasaharan Sethu,
Eliathamby Ambikairajah, David S. Taubman, and Julien Epps

Teaching Signal Processing Through Frequent and Diverse Design

A pedagogical approach

Digital Object Identifier 10.1109/MSP.2021.3057855
Date of current version: 28 April 2021

In this article, we describe and discuss the design-based approach for signal processing education at the undergraduate level at the University of New South Wales (UNSW) Sydney. The electrical engineering (EE) undergraduate curriculum at UNSW Sydney includes three dedicated signal processing courses as well as a design course that involves a major signal processing task. Design- and project-based teaching permeate the curriculum and form a primary method of content delivery to not only cement the understanding of theoretical concepts but also strengthen students' ability to apply these concepts to practical problems.

The incorporation of design-based teaching in the curriculum has been led by the signal processing group, with the pedagogical strategies evolving over nearly two decades of experience and innovation. The approach to project- and design-based teaching is based on the tenets of frequency, diversity, anchorage, significance, and autonomy. These tenets ensure that design projects are formulated in a way that conforms to Kolb's theory of experiential learning while simultaneously encompassing Bloom's taxonomy. Evidence is presented to demonstrate the effectiveness of the teaching strategy.

Introduction

In recent decades, digital signal processing (DSP) education has evolved into project- and problem-based learning, enabling students to engage with the theory in a hands-on format and develop the practical skills that lead to critical analysis and genuine comprehension of course content [1]–[4]. Although the benefits of design-based instruction are well established [5], [6], the actual formulation and construction, or "design," of these projects should be carefully considered to realize the pedagogical aims of the course.

Two pedagogical theories are of interest to us: Bloom's taxonomy of educational objectives [7] and Kolb's theory of experiential learning [8]. Bloom's taxonomy establishes a hierarchy of the learning stages that students can achieve. In this hierarchy, "knowing" a concept is the lowest level, while being able to "create" new work is the highest level; indeed, this is the ultimate goal of project-based learning. Kolb's experiential learning theory, on the other hand, provides a cycle of learning that a student can go through to construct knowledge from real experiences [9]. The four stages of this cycle are concrete experience, reflective observation, abstract conceptualization, and active experimentation [10]. Kolb's learning framework has been applied in engineering education for laboratory [11] and project [12] design.

Our approach to teaching at UNSW has evolved through nearly two decades of innovation in learning and teaching and is heavily based on the pedagogical theories mentioned. The frequency and diversity of the design tasks as well as their anchorage in DSP concepts ensure that students are given many opportunities, in a variety of contexts, to independently cement the theoretical knowledge they acquire in DSP courses. Our approach relies on the effective "design" of the design projects themselves. The emphasis on the tenets of frequency, diversity, anchorage, significance, and autonomy are a distinguishing feature of our teaching philosophy.

Signal processing education at UNSW

Modern undergraduate EE programs aim to produce industry-ready graduates [13] who possess strong theoretical knowledge as well as the necessary practical skills [4]. As such, graduates should have a deep understanding of the theoretical concepts and be capable of recruiting this understanding to analyze, create, and evaluate innovative solutions to real-world problems. A crucial aspect of the strategy to achieve this is to incorporate design and project-based learning into the engineering curricula.

A key element of the undergraduate signal processing education at UNSW Sydney is that design- and project-based subjects permeate the entire EE program [14]. Figure 1 depicts the signal processing stream and the complementary design stream within the EE program. Undergraduate EE students are introduced to design in their very first semester, with all students also required to take major design courses in the second, third, and fourth years in addition to the major research thesis in the fourth year. "Design proficiency" (DP) is a core focus of all learning [15].

As shown in Figure 1, the signal processing stream (depicted in blue) begins with a first-year subject on circuit theory, followed by a second-year circuits and systems subject that introduces continuous-time signal processing concepts. The delivery mode of this second-year course includes integrated tutorial–laboratory sessions, where students

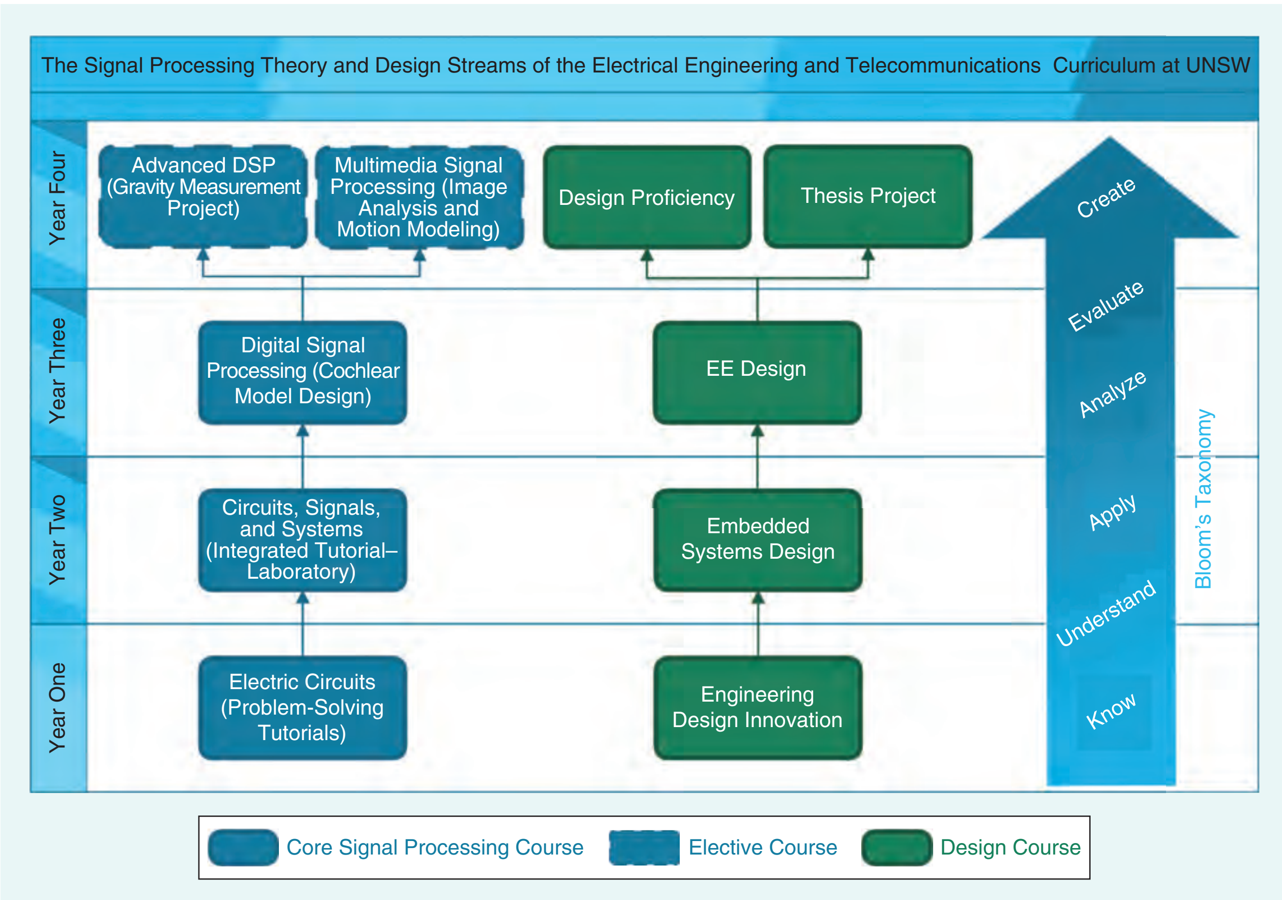

FIGURE 1. The structure, delivery modes, and philosophy of DSP and design subjects in the four-year bachelor of EE degree at UNSW Sydney. Solid borders denote core (compulsory) subjects, while dashed borders denote elective subjects, and arrows denote prerequisite dependencies. The theory subjects all include an experiential/project-oriented component, while the design subjects all include facets of signal processing theory and practice, helping to elevate student learning through Bloom's taxonomy.

work collaboratively on challenging problems that combine analytical, simulation, and practical work in a largely self-directed manner. These two subjects lay the foundation for the compulsory third-year subject, which introduces fundamental concepts in DSP. Two fourth-year electives, Advanced DSP (ADSP) and Multimedia Signal Processing (MMSP), build on the third-year knowledge and cover more advanced topics, such as adaptive and multirate signal processing, (introductory) maximum likelihood detection and estimation, and image and video processing.

In parallel with the signal processing stream, the EE program also includes a design stream comprising a set of design subjects that are compulsory in each year of study (shown in green in Figure 1). In the first year, students are challenged to develop unique solutions to simple design problems in a fun, competitive, team-based environment (for example, building a "Dalek" that rotates and points to a sound emitter) in teams of 4–8 students. Thus, they "learn by doing" basic concepts, such as basic signal detection and lowpass filtering.

The second-year embedded system design subject involves the use of modulation and demodulation, and so on, again in a problem-based learning mode. In the third year, pairs of students propose, design, and prototype electronic products they develop to solve practical problems. Often, this involves some DSP concepts they have learned in their third-year DSP course. An example project that was proposed by students is a vest equipped with ultrasonic emitters and microphones to detect obstacles to assist visually impaired people in navigating their surroundings. The vest would vibrate on the side of the obstacle, indicating its direction.

In their final undergraduate year, all EE students are required to pass a core DP subject [15], which has evolved into a demanding, yet rewarding, rite of passage that has drawn commendations from Engineers Australia (which administers Accreditation Board for Engineering and Technology-style accreditation) and industry employers alike. The DP course uses standards-based assessment, requiring all students to demonstrate strong analytical and practical DP in three challenging real-world tasks (for instance, controlling a motor in response to audio commands).

This unique DP subject also serves as an important quality assurance check on both the final-year graduating students themselves and on the courses that lead up to the final year. A third of the subject is focused on DSP design, and the learning expectations of this subject map to the very top of Bloom's taxonomy, meaning that students must engage the highest order of learning skills to take a series of industry-style design briefs all the way to fully functioning and fully explained engineering solutions that are each constructed and assessed live in the laboratory.

Philosophy

The philosophy that underpins our approach to signal processing education at UNSW is founded on project-based experiential learning. While traditional modes of teaching, such as lecture, tutorial, and lab-based course delivery, can achieve the required knowledge transfer, they can inadvertently leave many students at lower levels of Bloom's taxonomy, particularly in large classes with more than 100 students.

To address this, design-based teaching strategies are extensively employed to ensure that students have the opportunity to cement their knowledge in practical scenarios. Unlike traditional laboratory tasks, where each task is narrowly focused on a particular concept, design projects integrate many concepts; have multiple solutions; and challenge students to analyze the problem, evaluate various approaches, and create a suitable solution that satisfies the given requirements. Using the design-based methodology allows students to prioritize their time on concepts that they need to work on (instead of the instructor having to determine a suitable compromise for the entire class) and master each principle with a depth of knowledge and practice.

For nearly two decades, signal processing at UNSW Sydney has consistently been the discipline (within EE) that most strongly emphasizes experiential design. We employ design tasks with a mixture of individual and group-based learning so that every student develops a solid competency in DSP.

Bloom's taxonomy [7], shown in Figure 1, provides a hierarchical classification of the levels of learning that students can achieve. It comprises six levels, with the first three levels representing basic learning and the higher three levels covering deeper learning:

1) *Know*: This is the most basic level that can be achieved. Students learn concepts and can recall them without necessarily understanding them. For example, students may recall aliasing without understanding it and can state that avoiding it requires sampling above the Nyquist rate.
2) *Understand*: At this level, students can essentially restate concepts in their own words, which demands an understanding of the concept to be able to generate an equivalent definition. For example, students can explain *aliasing* in their own words rather than reciting the definition.
3) *Apply*: At this stage, students can apply the concepts they learned to problems they are given, like tutorial problems or basic laboratory tasks. For instance, given a signal and the sampling frequency, they can determine whether aliasing does or does not occur.
4) *Analyze*: This is effectively the first stage of deeper learning. At this stage, students master the concepts enough to be able to explain their parts or underlying mechanisms. For instance, at this stage, students can break down the phenomenon of aliasing to explain why it happens and which components of a signal would be affected by it.
5) *Evaluate*: This stage of Bloom's taxonomy is demonstrated by students being able to recruit their knowledge and understanding to evaluate propositions or approaches.

Effectively, they are able to make an assessment of statements or solutions. To this end, students need to be able to pick apart a concept to judge statements or solutions built upon it. In signal processing, this could be understood to mean that students are able to evaluate a solution against relevant or desirable criteria. Alternatively, when multiple solutions are available to them, they can evaluate and compare them. This is, in fact, a crucial ability that is required for the successful completion of open-ended design tasks.

6) *Synthesize*: This is the highest level of Bloom's taxonomy. Students are able to leverage their deeper understanding of the parts or underlying mechanisms of the concepts or phenomena and employ their ability to evaluate approaches to deliver effective solutions to new problems. That is, when given a problem like down-conversion of a signal, they can arrive at a solution that exploits aliasing by recruiting their deeper understanding of this concept.

While Bloom's taxonomy establishes the framework for students' learning, Kolb's system of experiential learning [8] provides a learning theory that can guide the delivery of knowledge to students. In DSP education, Kolb's cycle has been used in [16] to introduce experiential learning in the form of audio signal processing exercises that involve concrete experiences and active experimentation. In [17], Kolb's cycle was applied to the DSP stream of the curriculum to allow the program to instill skills, such as communications abilities and creative aptitudes, thereby increasing its appeal to students.

Kolb's cycle of experiential learning is an iterative process comprising four stages that form the ends of two axes, as shown in Figure 2: the vertical axis describes the continuum of grasping a concept either through concrete experience or abstract conceptualization, whereas the horizontal axis depicts the transformation of knowledge into understanding through reflective observation or active experimentation. The four stages in Kolb's cycle are explained here:

1) *Concrete experience*: In this stage, students encounter a new experience or concept and extract knowledge from that experience. As an illustration, we reuse the aliasing concept. Students may encounter the concept of aliasing in a simple experiment where the signal frequency is increased while the sampling frequency is kept constant. They are supervised and guided by the instructor, who plays the role of teacher, ensuring students are able to complete their experiment.
2) *Reflective observation*: At this point, students reflect on the experience and knowledge they gained and assimilate that knowledge. These reflections are articulated in discussions that are guided by the instructor, who focuses the session to consider the "why" and "how" underpinning their observations. For instance, reflecting on their observations of aliasing, students connect it to the Nyquist sampling theory by thinking about which sine waves may be represented by the time samples obtained through the sampling process.
3) *Abstract conceptualization*: Following reflection on the experience and assimilation of knowledge, students employ this knowledge to step to new ideas or new constructs. The instructor presents a challenge and asks questions that steer the thought process of students toward the particular concept or construct. For instance, students can then exploit their understanding of aliasing to figure out that one frequency may be shifted to another frequency by appropriate choice of the sampling frequency.
4) *Active experimentation*: Students then conduct new experiments to implement and verify the ideas resulting from their conceptualization of the knowledge gained. In the context of aliasing, students can then verify their proposed approach to frequency conversion. This experiment is supervised by the instructor, who engages with students to discuss with them their working and observations.

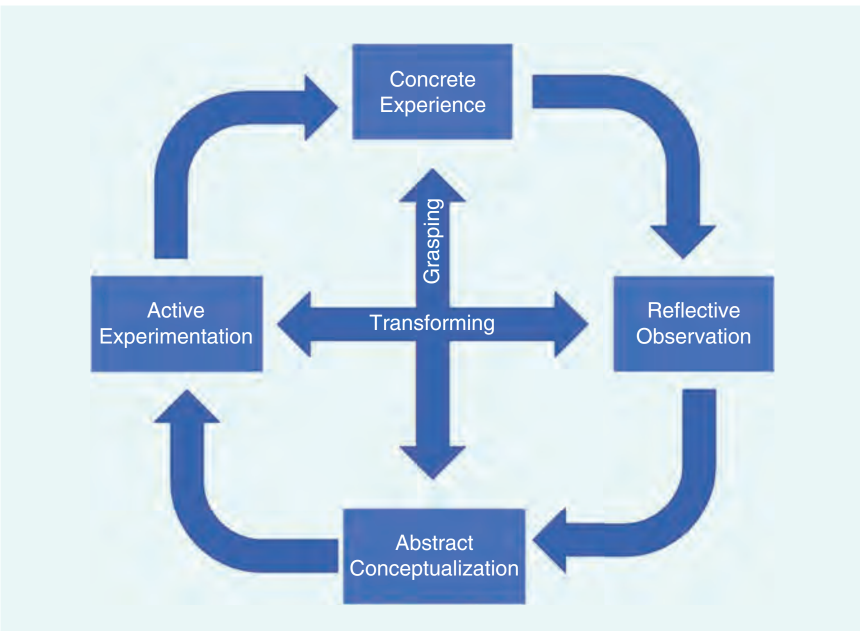

FIGURE 2. Kolb's experiential learning cycle showing the four stages of learning and the two axes connecting them [8].

Kolb's experiential learning and Bloom's taxonomy can be applied at various levels of the undergraduate program, ranging from the program level [18] to the subject [19] and individual component levels [11]. At the subject level, lectures, laboratory tasks, tutorials, and consultation sessions may be constructed to embody the experiential cycle, thus leading students through the four stages of the cycle described previously. For instance, lectures and tutorials can be used to introduce the concepts and facilitate the reflective observation part, whereas laboratory and consultation sessions allow students to experiment and conceptualize their knowledge. At the component level, Kolb's theory may be used as a framework for laboratory tasks or design projects such that activities that facilitate each of the four stages are incorporated into the task or project.

Bloom's and Kolb's systems of learning can be brought together into a coherent framework [12], [18], [20] to achieve effective and deep learning. As a flexible system for experiential learning, Kolb's theory can be used to develop a teaching strategy that would facilitate the movement of students up the levels of Bloom's taxonomy. In this way, every time an instructor guides a student through the complete Kolb's cycle, the student's movement from active experimentation to abstract conceptualization aids in achieving a progressively deeper understanding of the concepts involved.

Thus, we view each repetition of Kolb's cycle as part of a spiral path, leading students up these levels (of Bloom's taxonomy), and our approach to the delivery of signal processing knowledge, which is described in detail in the next section, embodies this combined framework. To take full advantage of this pedagogical framework, our philosophy regarding experiential learning has evolved to employ the following tenets:

1) *Frequency*: To reinforce the knowledge gained and consolidate their practical skills, we continually and persistently expose students to experiential learning. As can be seen from Figure 1, the design tasks featuring signal processing that our students undertake are stacked over the degree with multiple projects in each year. The frequency of these opportunities ensures that students go through many iterations of Kolb's cycle.
2) *Diversity*: Students are diverse and so are their interests and learning styles. Catering to these different interests and styles allows us to better motivate and engage students [4], [21]. Thus, the design tasks employ a variety of contexts and formulations across the signal processing and design streams illustrated in Figure 1. The wide range of design projects provides multiple perspectives on signal processing concepts, which increases the likelihood of students grasping the concepts and deepening their understanding.
3) *Anchorage*: The design tasks are firmly anchored in real applications. They tap into practical or even familiar contexts, such estimating the signal parameters in a power system. Also, they often take advantage of the research labs, like the speech, multimedia, and radar signal processing labs. This provides for an improved concrete experience and helps students see the relevance of the concepts they are learning. Thus, the familiarity of the context or application enhances their motivation to engage in the learning process and improves students' understanding of these concepts. Examples of this are the cochlear signal processing and gravity measurement projects that are described in the "Example Design Projects" section.
4) *Significance*: This quality concerns the depth and difficulty of the design tasks. Specifically, the objectives of the tasks are formulated to be significant and challenging, with a variety of possible solutions. Thus, students are required to exercise their analytical and decision-making skills to devise a suitable solution. In the cochlear signal processing project described in the "Example Design Projects", students are required to bring together many of the concepts they learned to complete the project. In the gravity measurement exercise, students are expected to model the system, that is, to relate the observed frequency to gravity through the equations of motion and the Doppler shift. After the modeling, they proceed to employ advanced signal processing concepts to develop a solution. Significance goes hand in hand with the final tenet of autonomy.
5) *Autonomy*: Students are given significant autonomy in facing design challenges. Although support is provided, answers are not given. The instructors play the role of mentors, challenging and guiding students to think critically to arrive at the answers themselves. They are required to analyze the problem statement, derive the requirements, evaluate solutions, and exercise their judgment. Autonomy has been found to enhance motivation, improve engagement, and increase comprehension [22]. It also helps students develop their ability to work independently, which is a necessary quality in the professional environment.

Approach

The DSP stream currently comprises three DSP subjects and two lower-level supporting subjects in year one and year two. The design pillar, on the other hand, is a progression of one design subject per year, culminating in the DP subject and the final-year thesis. All of the DSP subjects involve substantial design tasks/projects, and all of the design subjects, and in particular DP, include significant DSP components. Consequently, there is a diverse range of DSP design tasks that permeate the entire program and are frequently encountered by students.

Although the subjects and design projects are underpinned by the experiential learning pedagogical philosophy, their exact formulation varies among subjects to account for the increasing depth of understanding that evolves along the engineering degree. In this section, we describe the structure of these subjects and highlight their relationships to the underlying philosophy.

In the first two years, students learn basic circuit theory (year one) and linear systems theory (year two). The third-year foundational subject on DSP is a core subject that is taken by all undergraduate students in the program and introduces them to the fundamentals of DSP. The topics covered include sampling and aliasing, discrete versions of Fourier analyses, linear discrete-time system theory, z transforms, finite-impulse response and infinite impulse response filter designs, and introductory multirate processing.

This DSP subject incorporates a number of "learning and teaching" elements designed in accordance with Kolb's framework to take students up the levels of Bloom's taxonomy. The lectures explain new ideas to students and take them up the first two levels of Bloom's taxonomy. These are then supported by two components that implement the experiential learning cycle:

- Tutorials and laboratories are integrated in a single tutorial–lab unit and involve problems that are designed such that the analytical solutions can be implemented as programs in a numerical computing environment such as MATLAB. This approach fits quite well with Kolb's cycle, as illustrated in Figure 3. Students first apply their understanding of the new ideas they encountered in the lectures to solve the tutorial problems and implement them in the laboratory (the third step on Bloom's taxonomy). This affords students the opportunity to visualize and understand the nuances of their analytical work (the fourth step on Bloom's taxonomy) and to validate it (the fifth step). Students are allowed to complete these tutorial–lab problems at more or less their own pace with weekly tutor support time-tabled as part of the course. As an example, a tutorial–lab problem might require students to design a filter via pole-zero placement, derive the transfer function, sketch the magnitude response, and then implement the filter in MATLAB and validate its properties. The tight feedback loop involving validation via MATLAB simulations and the opportunity to redo the problems allow students to iterate through Kolb's framework multiple times, if necessary, to progress up the fourth and fifth and, to a lesser extent, the sixth steps of Bloom's taxonomy.
- Students are also assigned to work on a term-long project that is more open ended and requires them to undertake a specific task but with the freedom to choose from among various paths to reach their goal. This lets students evaluate multiple options; compare among them; and, finally, synthesize a significant, coherent body of work applying most of the concepts introduced in this subject. The project enables students to experiment, analyze, evaluate, and consult with the tutors and course lecturers to finally arrive at their solution. This process iterates through Kolb's framework, allowing students to reach the top of Bloom's taxonomy. The project is anchored in real-world applications, such as cochlear signal processing, to give students a feel for DSP concepts that might be used in a real-world or complex system. The cochlear signal processing design project used in 2018 and 2019 is described in the "Example Design Projects" section.

The fourth-year elective ADSP subject builds on the material learned in the third-year course. In addition to introducing mathematical rigor that underpins the basic concepts of DSP, this subject advances students' knowledge, exposing them to statistical signal processing, detection and estimation theory, adaptive filters, and time-frequency representations. In addition to the lectures, the subject makes use of carefully designed tutorial questions and challenge problems and includes a laboratory program that comprises small laboratory tasks leading to a significant design project.

As this is an advanced subject, the lectures, tutorials, and labs are designed to achieve the highest level of Bloom's taxonomy. The course design reflects this, placing strong emphasis on anchorage, significance, and autonomy (particularly for the final project). Group work and collaboration are strongly encouraged in most aspects of the course, and the laboratory program requires students to regularly demonstrate their competency at the "evaluate" and "create" levels of Bloom's taxonomy.

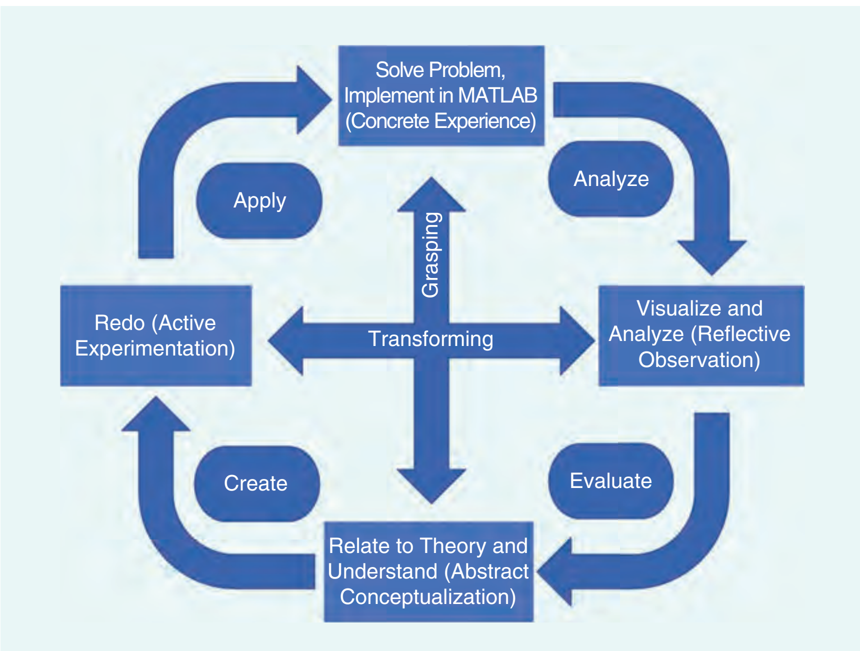

FIGURE 3. Kolb's cycle representation of the integrated tutorial–laboratory approach of the DSP subject showing the corresponding levels of Bloom's taxonomy.

Small design tasks allow students to evolve the knowledge they gained in lectures by deepening and expanding their understanding and advancing some new concepts not covered in lectures. For example, while lectures on adaptive filtering cover the Wiener filter, the associated lab task goes further, explaining the minimum variance distortionless response filter. These small tasks then lead to a major design project in the second half of term, and students are required to submit a report at its conclusion.

The entire laboratory program provides breadth of coverage and diversity as well as depth and significance, and the project, in particular, embodies multiple iterations of Kolb's cycle. Projects are often anchored in the research area of the course lecturer, with example projects that have been used including the analysis of bat echolocation calls, measurement of power system parameters, and gravity measurement project (described in the "Example Design Projects" section).

The elective MMSP subject focuses on the extension and application of signal processing concepts into two and three dimensions (images and video), with an emphasis on understanding concepts specific to multimedia signals, e.g., shape, orientation, color, motion, formats, and representation schemes. The subject places strong emphasis on the practical implementation of the associated algorithms in real programming environments (C/C++). This is a deliberate choice, made to give students confidence in the programming environments they will employ after they graduate.

While the subject includes significant fundamental material of a theoretical nature, it also incorporates a substantial laboratory program to provide students with a balanced and efficient path up Bloom's taxonomy. The practical component includes projects where students write code in C/C++, starting only from the theory. In this manner, they acquire coding, debugging, and memory management skills as well as achieving deeper understanding while experiencing the design perspective (no single correct solution). Example projects are

- image up-sampling and down-sampling using disciplined interpolation kernels to avoid aliasing
- image skewing and rotation with asymptotically reversible filters
- practical bilevel and gray-scale morphology
- global motion estimation based on robust keypoint detection
- generation of texture maps through robust local power spectrum analysis.

Finally, Electrical DP is a final-year core subject that is taken by all EE students at UNSW [15]. It is a somewhat unique course, with the primary aim to test students' proficiency at EE design over a set of three challenges. In contrast to the final-year thesis, which spans the entire final year, DP tasks must be undertaken and assessed within 4-h laboratory slots, subject to constraints on the choice of components and tools.

One of the three design tasks, which covers three laboratory slots, is focused on DSP and involves a design challenge where students have to develop a solution to meet a given set of requirements, prototype it, validate the design, and demonstrate that it works. Students operate with a very high level of autonomy in completing the tasks and, to do well in the course, they must demonstrate the ability to operate at the highest level of Bloom's taxonomy.

The DP subject is entirely practice based and, therefore, most completely encompasses a comprehensive pedagogical framework, integrating Bloom's taxonomy with Kolb's experiential theory. Each task constitutes a significant challenge, requiring students to iterate through the experiential learning cycle to arrive at a solution they can demonstrate. Thus, analysis, evaluation, and experimentation as well as research are essential activities. An example project is to design and implement a system that allows a light bulb to be reliably turned on and off by clapping. As it challenges students to apply and expand their knowledge, this course also serves as validation that previous courses in DSP in the program have been effective at bringing them to this level.

Having outlined our design-based approach, it is important to add a remark on the workload involved. It is well known that benefits of design and project-based learning come at an added cost in terms of resources and staff workload. Although the details of the workload and resources requirements vary among subjects, they are nonetheless manageable, as evidenced by our successful implementations for cohorts ranging from around 60 students in ADSP and MMSP to 150 students in DP and more than 200 students in the third-year DSP subject. Our experience shows that, in addition to access to high-quality laboratory and hardware facilities, successful delivery of our philosophy requires the provision of adequate support in terms of consultation periods, which is effectively handled through a modest increase in staff commitment.

Example design projects

Recruiting current technologies and applications into the teaching strategy provides a fresh perspective on "old" signal processing concepts [23]. Thus, design tasks are anchored in practical applications, such as speech, image and video, biomedical, power systems, radar, and array sensors. In this section, we briefly describe and discuss two example projects.

Cochlear signal processing

This project is motivated by the observation that the first DSP course tends to be a significant jump in abstraction for most students, which is a consequence of them being introduced to a large number of mathematical concepts with insufficient time to consolidate the ideas with practical examples. This design project is crafted as a platform that brings together almost all of the fundamental DSP concepts taught in an introductory course and lends itself as a suitable candidate for project-based learning [24].

The goal of the cochlear signal processing project was twofold: 1) to involve students in a single project that can reinforce most of the theoretical concepts introduced in the course and 2) to provide them with a real-world example that engages their interest and gives them signal processing ideas they can continue to think about beyond the course. In an introductory lecture, we provide students with basic information on the physiology of hearing:

- introduction to the human auditory system
- operation of the outer ear and middle ear as a combined bandpass filter
- operation of the inner ear as a cascade of filters that perform spectral analysis
- operation of the inner hair cells in signal encoding at a reduced sampling rate.

Then, the entire project involves implementing

- a combined digital filter model of the outer and middle ear
- a cascaded filter (transmission line) model of the basilar membrane, acting as a spectrum analyzer
- a nonlinear model of the inner hair cells.

Table 1. The DSP concepts that each part of the cochlear project will help consolidate.

Task	DSP Concepts Covered
Outer and middle ear model	Pole-zero placement on the *z*-plane Transfer functions and *z* transform Magnitude response estimation Filter structure
Inner ear model	Impulse response Resonant filters and selectivity Impulse invariant and bilinear transforms Digital filter design and implementation Discrete-time Fourier transforms Stability Characterization of magnitude response Spectral analyses
Inner hair cell model	Rectification and envelope estimation Down-sampling

Table 2. The DSP concepts that each part of the gravity measurement project will help consolidate.

Task	DSP Concepts Covered
Signal modeling: model the sound of the siren and relate to acceleration due to gravity	Signal representation Sampling
Filtering: filter background noise to enhance the siren signal-to-noise ratio	Digital filter design Adaptive versus nonadaptive filters Digital filter implementation Discrete-time Fourier transforms
Signal representation: transform the signal to the frequency or time-frequency domain	Time-frequency representations The short-time Fourier transform Real versus complex signals Spectral leakage
Parameter estimation: estimate the siren frequency versus time and then obtain an estimate of the acceleration due to gravity	Maximum likelihood estimation Practical estimators Least squares estimation

Table 3. A summary of the responses to feedback questions about the cochlear signal processing project.

Question	Percentage (Agree + Strongly Agree)
1) As a result of working on this project, my understanding of DSP concepts improved.	76
2) As a result of working on this project, I have a better understanding of how analog systems may be modeled digitally.	87
3) The project was challenging and encouraged me to learn.	78
4) This project provided me with an opportunity for group and collaborative learning.	71
5) As a result of working on this project, my proficiency with MATLAB improved.	85

These three aspects cover almost all of the DSP concepts taught in the course, as outlined in Table 1.

Gravity measurement project

In this project, students record a falling siren and employ signal processing techniques to obtain a measurement of the acceleration due to gravity [25]. The project uses the principle of “unfamiliar concepts in familiar contexts” to cast the measurement of gravity, which is familiar even to high school physics students, into a signal processing challenge that builds on the concepts of filtering, time-frequency analysis, detection, and estimation.

The project starts with an experimental stage to collect data by dropping a monotone siren from a stationary position at some height, h, directly above a microphone and recording its sound before and during the fall. Multiple recordings are obtained under various conditions, including against a quiet background and in noisy settings. Students are then required to develop a solution to process the recorded files to filter out background noise, estimate the siren frequency at rest and during the fall, and then fit a line to the changing frequency to finally calculate the acceleration due to gravity from the observed Doppler shift. A list of the signal processing concepts that are covered by this project is given in Table 2.

Evidence

The design-based experiential learning approach over the entire EE degree, and especially within signal processing, has produced cohorts of students who have a greater depth of knowledge and applied skills that are pertinent for their careers in industry. This teaching philosophy has evolved through many years of continual evaluation and improvement of the curriculum in general and the DSP stream in particular.

Evidence for the effectiveness of our experiential approach to learning and teaching signal processing includes indicators that are both qualitative and quantitative. Indicators of a quantitative nature include survey results from the DSP and ADSP subjects that ran the two projects described in the previous section, student performance in the DP subject, and student retention rate in the advanced signal processing courses.

Students who undertook the cochlear signal processing project in the third-year introductory DSP course in 2018 were surveyed at the end of the term (prior to their final exam) about their learning experience in the project. A number of questions, covering both the perceived quality and effectiveness of the project, were asked, and almost 90% of the enrolled students responded to the survey (179 out of 205). The response to each question could be one of five options: strongly disagree, disagree, neither agree nor disagree, agree, and strongly agree.

A summary of the results giving the percentage of agree and strongly agree responses is provided in Table 3. We see that more than three quarters of students reported that their understanding of DSP had improved (Q1), whereas 87% responded that the project gave them a better understanding of the digital modeling of analog systems, which is precisely

Table 4. The topics used in the self-benchmarking test in ADSP.

Topic Number	Topic	Topic Number	Topic
1	Sampling (definition and mathematical expressions)	11	Relationships among the z, Laplace, and Fourier representations
2	Aliasing (frequency domain and time domain)	12	Filters (difference equation, transfer function, and frequency response)
3	Spectra (definition, meaning, and properties)	13	Filter structures (direct form and canonical form)
4	Convolution (meaning, definition, and operations)	14	Filter properties (gain, phase, and group delay)
5	Relationship between discrete and continuous domains	15	Filters types (linear/minimum/maximum phase and all pass)
6	Linearity and time invariance (meaning, definition, and so on)	16	Manipulating the transfer function (poles and zeros, factorization, partial fractions)
7	Discrete-time (linear time-invariant) systems	17	Fourier transforms and series
8	Stability (bounded-input, bounded-output, and others)	18	Relationships among signal representations
9	The frequency/impulse responses and their relationships	19	Filter design methods
10	The z transform (need, meaning, properties, and operations)	20	Quantization and rounding effects

due to the anchorage of the project in the practical application. Overall, it is observed that the students valued the experience and benefited from it. Similar feedback was also obtained in 2019, when the project was run again but with the transmission line replaced by a parallel filter bank model.

A self-benchmarking test was given to ADSP students both at the start and end of the term. Students were asked at the start of term one, 2020, to rate their understanding of the fundamental topics listed in Table 4 on a scale of zero to 10, with zero indicating that they have not encountered the concept before and one to 10 (10 being the highest level) giving a rating of their understanding of the topic. The same self-benchmarking test was given to the cohort at the end of the term, and the responses were compared to ascertain the progression in students' self-assessed understanding. For the start-of-term test, we had responses from 33 students, whereas at the end of the term, 16 responded. Only the responses of students who took both tests were used in the analysis. The mean scores for each topic as well as the changes in mean ratings are plotted in Figure 4. It is clear that students felt that they benefited from their learning experience and had better command of all of the topics.

The cochlear survey and self-benchmarking test are two measures that rely on students' feedback. We now present two additional measures that are not based on students' perceptions. These serve to benchmark the previous two indicators and to affirm our conclusions regarding the effectiveness of the learning and teaching philosophy.

The first of these measures is based on student performance in the DP subject, which was described in the "Approach" section. The DP subject was developed and initially taught by the signal processing group. It tests students on their ability to apply their knowledge to develop and demonstrate solutions to challenging design tasks. In this way, it requires that students show a high level of command of the concepts they learned in the core subjects and the capacity to manipulate these concepts to complete the tasks.

As systems and control education at UNSW has employed a more traditional approach, the control design task provides a useful benchmark against which we may gauge the effectiveness of our approach to students' learning. A comparison of the marks achieved by students in the signal processing and control tasks is presented in Figure 5. The graph shows the 25th and 75th percentile bands for the marks achieved in the signal processing and control design tasks over the years 2011 to 2018. The results show that students performed better in the signal processing tasks over all of the years shown. Note that the yearly cohort sizes for DP tend to vary between 120 and 180.

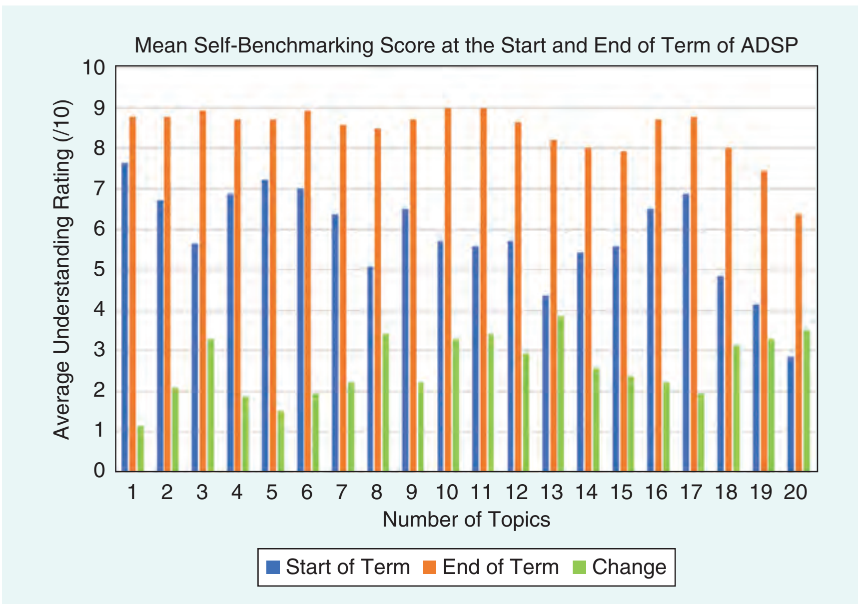

FIGURE 4. The results of the self-benchmarking test administered in the ADSP subject: the mean score at the start and end of the term.

Another quantitative measure that points to the effectiveness of our approach in engaging and motivating students is the retention rate from the core third-year DSP subject to the fourth-year elective ADSP and MMSP courses. The number of enrollments in each course for the years 2017 to 2019 as well as their fraction of the number of students in the DSP course in the previous year are shown in Figure 6. We see an increasing trend, with more than 30% of the cohort choosing to take the fourth-year DSP courses. This is remarkable, given that more than 20 electives are on offer to fourth-year EE students.

A variety of qualitative evidence also supports the effectiveness of our approach. This includes increased take-up of signal processing theses in the final year, growing interest in extracurricular projects (especially in the various signal processing research labs in the school), and the number of students choosing to pursue a Ph.D. degree in signal processing both within and outside UNSW. However, perhaps one of the most pertinent pieces of evidence comes from the fact that our signal processing students have won major international competitions. In fact, undergraduate teams have participated in five of the seven IEEE Signal Processing Cup international competitions, consistently performing very well on the international stage, including winning first place in 2017 [26] and second place in 2019 [27].

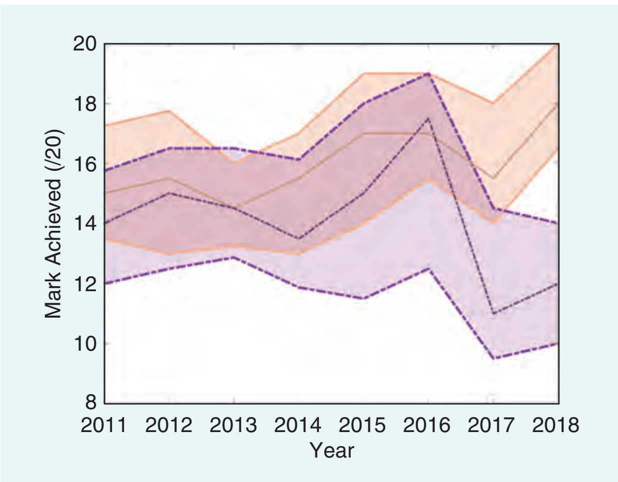

FIGURE 5. The student performance in the signal processing and systems and control design tasks in the DP subject over the years from 2011 to 2018. The bottom and top solid lines are the 25th and 75th percentiles for the signal processing tasks, while the middle solid line is the median. Similarly, the dashed lines pertain to the systems and control tasks.

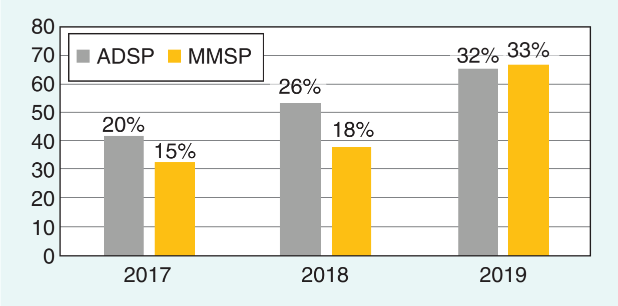

FIGURE 6. The student participation in the fourth-year elective signal processing courses. These results reflect the retention rate from the compulsory third-year DSP course.

Conclusions and future outlook

Signal processing education has been evolving to incorporate more design-based learning. At UNSW, our approach to DSP education encompasses both Bloom's taxonomy and Kolb's theory of experiential learning and exposes students to a continuous and varied stream of design projects. This approach is based on the five tenets of frequency, diversity, anchorage, significance, and autonomy. These tenets serve to ensure that students are challenged frequently and meaningfully to be able to iterate through Kolb's cycle to efficiently climb to the top of Bloom's taxonomy.

As an evolving strategy, our approach is continually being improved in response to sharpened evaluation processes and student feedback. For instance, plans are already in place to employ the self-benchmarking test in the third-year DSP course to gauge the effectiveness of the subject in delivering the various concepts to students. The results can then be correlated with the survey results for the project in DSP as well as with the self-benchmarking test in ADSP.

Authors

Elias Aboutanios (elias@ieee.org) received his bachelor's degree in electrical engineering in 1997 from the University of New South Wales (UNSW), Australia, and his Ph.D. degree in 2003 from the University of Technology Sydney, Australia. He is an associate professor at the School of Electrical Engineering and Telecommunications, UNSW, Sydney, 2052, Australia. From 2003 to 2007, he was a research fellow with the Institute for Digital Communications, University of Edinburgh. He is an associate editor of *IEEE Transactions on Signal Processing* and *IET Signal Processing*. He is a member of the IEEE Sensor and Multichannel Technical Committee. He established and led the UNSW EC0 CubeSat Project, which culminated in the launch of the satellite in 2017. His research interests include statistical signal processing, particularly signal detection and parameter estimation, for applications such as radar, the global navigation satellite system, smart grids, and nuclear magnetic resonance spectroscopy. He is a Senior Member of IEEE.

Vidhyasaharan Sethu (v.sethu@unsw.edu.au) received his B.E. degree from Anna University, Chennai, India, in 2005, his M.Eng.Sc. degree in signal processing, and his Ph.D. degree in speech signal processing from the University of New South Wales (UNSW), Sydney, Australia, in 2006 and 2010, respectively. He is a senior lecturer at the School of Electrical Engineering and Telecommunications, UNSW, Sydney, 2052, Australia. From 2010 to 2013, he was a postdoctoral fellow at the Speech Processing Research Group, UNSW. He has coauthored approximately 90 publications and serves on the editorial board of *Computer Speech and Language*. His research interests include the application of

machine learning to speech processing and affective computing, speaker recognition, and computational paralinguistics. He is a Member of IEEE.

Eliathamby Ambikairajah (e.ambikairajah@unsw.edu.au) received his B.Sc. (Eng.) (Hons.) degree from the University of Sri Lanka and his Ph.D. degree from Keele University, United Kingdom. Since 2008, he has been a professor with the School of Electrical Engineering and Telecommunications at the University of New South Wales, Sydney, 2052, Australia. He was appointed as head of electronic engineering and, later, dean of engineering at the Athlone Institute of Technology in the Republic of Ireland from 1982 to 1999. He was an associate editor of *IEEE Transactions on Education* from 2012 to 2019. His research interests include speaker and language recognition, emotion detection, and cochlear modeling. He is a Senior Member of IEEE.

David S. Taubman (d.taubman@unsw.edu.au) is with the School of Electrical Engineering and Telecommunications at the University of New South Wales, Sydney, 2052, Australia. He is also the director and founder of Kakadu Software Pty. Ltd. Prior to these roles, he spent four years at Hewlett-Packard's research laboratories in Palo Alto, California. He contributed extensively to the JPEG2000 suite of standards, developing the core coding algorithms for both parts 1 and 15 (HTJ2K) as well as JPIP (part 9) for interactive communication. He has also made many contributions to scalable video coding. He is a coauthor with Michael Marcellin of the book *JPEG2000: Image Compression Fundamentals, Standards and Practice* and recipient of two IEEE Society Best Paper awards. He is a Fellow of IEEE.

Julien Epps (j.epps@unsw.edu.au) received his B.E. and Ph.D. degrees from the University of New South Wales, Sydney (UNSW), Australia. Following appointments as a senior research engineer with Motorola Labs and as a senior researcher with National ICT Australia, he joined the School of Electrical Engineering and Telecommunications at UNSW, Sydney, 2052, Australia, in 2007, where he is now a professor and head of the school. He serves as an associate editor of *IEEE Transactions on Affective Computing.* His current research interests include characterization, modeling, and classification of mental state from behavioral signals, such as speech, eye activity, and head movement. He is a Member of IEEE.

References

[1] E. Ambikairajah and J. Epps, "Project-based learning in digital signal processing: Development and experiences," in *Proc. 2011 Digital Signal Process. Signal Process. Educ. Meeting (DSP/SPE)*, pp. 506–511. doi: 10.1109/DSP-SPE.2011.5739266.

[2] E. Ambikairajah, T. Thiruvaran, and R. Eaton, "Enhancement of traditional mode of teaching based on the experience of project-based learning," in *Proc. 2013 IEEE 8th Int. Conf. Ind. Inf. Syst.*, pp. 665–670. doi: 10.1109/ICIInfS.2013.6732064.

[3] A. Pardo, "Problem-based learning combined with project-based learning: A pilot application in digital signal processing," in *Proc. 2014 11th Tecnologias Aplicadas a la Ensenanza de la Electronica (Technologies Applied to Electronics Teaching) (TAEE)*, pp. 1–5. doi: 10.1109/TAEE.2014.6900128.

[4] H. Godrich, A. Nehorai, A. Tajer, M. S. Greco, and C. Zhang, "Special article series on signal processing education via hands-on and design projects [from the guest editors]," *IEEE Signal Process. Mag.*, vol. 34, no. 1, pp. 13–15, 2017. doi: 10.1109/MSP.2016.2620199.

[5] J. W. Thomas, *A Review of Research on Project-Based Learning*. San Rafael, CA: Autodesk Foundation, 2000.

[6] J. Macias-Guarasa, J. M. Montero, R. San-Segundo, A. Araujo, and O. Nieto-Taladriz, "A project-based learning approach to design electronic systems curricula," *IEEE Trans. Edu.*, vol. 49, no. 3, pp. 389–397, 2006. doi: 10.1109/TE.2006.879784.

[7] B. S. Bloom, P. Airasian, K. Cruikshank, R. Mayer, P. Pintrich, J. Raths, and M. Wittrock, *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. New York: Longman, 2001.

[8] D. A. Kolb, *Experiential Learning: Experience as the Source of Learning and Development*. Upper Saddle River, NJ: Pearson FT Press, 2014.

[9] S. Yardley, P. W. Teunissen, and T. Dornan, "Experiential learning: Transforming theory into practice," *Med. Teach.*, vol. 34, no. 2, pp. 161–164, 2012. doi: 10.3109/0142159X.2012.643264.

[10] J. Eyler, "The power of experiential education," *Liberal Educ.*, vol. 95, no. 4, pp. 24–31, 2009.

[11] M. Abdulwahed and Z. K. Nagy, "Applying Kolb's experiential learning cycle for laboratory education," *J. Eng. Educ.*, vol. 98, no. 3, pp. 283–294, 2009. doi: 10.1002/j.2168-9830.2009.tb01025.x.

[12] L. Schatzberg, "Applying Bloom's and Kolb's theories to teaching systems analysis and design," in *Proc. Inf. Syst. Educ. Conf.*, 2002, vol. 19, pp. 1–8.

[13] J. J. Zacharias and J. M. Conrad, "A survey of digital signal processing education," in *Proc. 2007 IEEE SoutheastCon*, pp. 322–327. doi: 10.1109/SECON.2007.342910.

[14] "Electrical engineering (honours)." UNSW Sydney. Sept. 2020. https://www.engineering.unsw.edu.au/study-with-us/undergraduate-degrees/electrical-engineering-honours-0 (accessed Feb. 12, 2021).

[15] E. Aboutanios and D. Taubman, "Design proficiency: A novel Electrical Engineering course," in *Proc. IEEE Int. Conf. Teaching, Assessment Learn. Eng. (TALE),* Hong Kong, Aug. 2012, pp. W1D–17–W1D–21. doi: 10.1109/TALE.2012.6360289.

[16] J. W. Pierre, R. F. Kubichek, and J. C. Hamann, "Reinforcing the understanding of signal processing concepts using audio exercises," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Process.*, 1999, vol. 6, pp. 3577–3580. doi: 10.1109/ICASSP.1999.757616.

[17] G. Plett, R. Ziemer, M. Ciletti, R. Dandapani, T. Kalkur, and M. Wickert, "Experiences in updating the ECE curriculum with signal processing first and Kolb/4 MAT pedagogy," in *Proc. 2006 Annu. Conf. Expo.*, pp. 11–614.

[18] J. N. Harb, S. O. Durrant, and R. E. Terry, "Use of the Kolb learning cycle and the 4MAT system in engineering education," *J. Eng. Educ.*, vol. 82, no. 2, pp. 70–77, 1993. doi: 10.1002/j.2168-9830.1993.tb00079.x.

[19] R. A. Howard, C. A. Carver, and W. D. Lane, "Felder's learning styles, Bloom's taxonomy, and the Kolb learning cycle: Tying it all together in the CS2 course," in *Proc. 27th SIGCSE Tech. Symp. Comput. Sci. Educ.*, 1996, pp. 227–231.

[20] E. J. Murphy, "Prior learning assessment: A review of Bloom's taxonomy and Kolb's theory of experiential learning: Practical uses for prior learning assessment," *J. Continuing Higher Educ.*, vol. 55, no. 3, pp. 64–66, 2007. doi: 10.1080/07377366.2007.10400135.

[21] T. Schäck, M. Muma, and A. M. Zoubir, "Signal processing projects at Technische Universität Darmstadt: How to engage undergraduate students in signal processing practice," *IEEE Signal Process. Mag.*, vol. 34, no. 1, pp. 16–30, 2017. doi: 10.1109/MSP.2016.2619218.

[22] M. C. Pérez, Á. Hernández, and R. Mateos, "Promoting student autonomy in the electronic design course," in *Proc. Design Circuits Integrated Syst.*, 2014, pp. 1–6.

[23] S. Kanna, W. von Rosenberg, V. Goverdovsky, A. G. Constantinides, and D. P. Mandic, "Bringing wearable sensors into the classroom: A participatory approach [SP Education]," *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 110–130, 2018. doi: 10.1109/MSP.2018.2806418.

[24] E. Ambikairajah and V. Sethu, "Cochlear signal processing: A platform for learning the fundamentals of digital signal processing," in *Proc. IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, 2020, pp. 9229–9233. doi: 10.1109/ICASSP40776.2020.9054297.

[25] E. Aboutanios, "Bringing signal processing to high school students—A novel gravity measurement experiment," in *Proc. IEEE Int. Conf. Teaching, Assessment Learn. Eng. (TALE),* Kuta, Aug. 2013, pp. 419–422. doi: 10.1109/TALE.2013.6654474.

[26] C. T. Jin, M. Davies, and P. Campisi, "Embedded systems feel the beat in New Orleans: Highlights from the IEEE Signal Processing Cup 2017 student competition [SP Competitions]," *IEEE Signal Process. Mag.*, vol. 34, no. 4, pp. 143–170, July 2017. doi: 10.1109/MSP.2017.2698075.

[27] A. Deleforge, D. Di Carlo, M. Strauss, R. Serizel, and L. Marcenaro, "Audio-based search and rescue with a drone: Highlights from the IEEE signal processing cup 2019 student competition [SP Competitions]," *IEEE Signal Process. Mag.*, vol. 36, no. 5, pp. 138–144. doi: 10.1109/MSP.2019.2924687.

SP

Felipe Tobar, Felipe Bravo-Marquez, Jocelyn Dunstan,
Joaquín Fontbona, Alejandro Maass,
Daniel Remenik, and Jorge F. Silva

Data Science for Engineers

A teaching ecosystem

©SHUTTERSTOCK.COM/LUCKYSTEP

We describe an ecosystem for teaching data science (DS) to engineers that blends theory, methods, and applications, developed at the Faculty of Physical and Mathematical Sciences (FCFM is its Spanish acronym), Universidad de Chile, over the last three years. This initiative has been motivated by the increasing demand for DS qualifications both from academic and professional environments.

The ecosystem is distributed in a collaborative fashion across three departments in FCFM and includes postgraduate programs, courses, professional diplomas, data repositories, laboratories, trainee programs, and internships. By sharing our teaching principles and the innovative components of our approach to teaching DS, we hope our experience can be useful to those developing their own DS programs and ecosystems. The open challenges and future plans for our ecosystem are also discussed at the end of the article.

Introduction

Interdisciplinarity to embrace new challenges

The taxonomy of academic branches reflects the necessities of a society in space and time and is, therefore, subject to both gradual and sudden changes, just as the evolution of science [1]. Early universities dealt with subjects such as theology and natural philosophy; wars catalyzed the teaching of engineering (both civil and military), while schools in rural areas have grown to focus on agricultural studies, and business schools have arisen near financial districts. This suggests that the division we impose over different branches of knowledge and, in particular, of science is, to a large extent, purely instrumental: it obeys our own necessities and not an evident or natural segmentation [2].

As the necessities and interests of societies change over time, academic branches evolve both in depth and scope. This problem-driven reformulation promotes the creation of new relevant fields and academic branches. However, this does not always occur in a timely manner but is, instead, a lengthy process whose timing often lags behind urgent societal demands.

Digital Object Identifier 10.1109/MSP.2021.3053551
Date of current version: 28 April 2021

In response to increasing complex practical needs and societal demands, approaches that rely on the interaction of existing well-studied branches of knowledge have recently come into focus. This interaction between or among disciplines is what is referred to as *interdisciplinarity/multidisciplinarity* [3] and has become essential to address current challenges effectively. Multidisciplinary/interdisciplinary approaches are effective because the skills and abilities required to confront today's challenges are segregated across different disciplines due to the outdated taxonomy of knowledge that was once imposed under different conditions.

The case of DS

A contemporary instance of this phenomenon is DS [4], [5]. From the public sector, industry, and academia, a number of agents are demanding (and also offering) solutions that are labeled as DS or related terms, such as artificial intelligence (AI), machine learning (ML), data mining, and big data.

Briefly put, a DS task is one that involves some of the following stages: acquisition, curation, transmission, processing, analysis, interpretation, and visualization of some form of information content. The skills needed to address current challenges in DS are mainly those found in machine learning [6], mathematics (optimization [7], probability and statistics [8]), computer science (data mining, semantic web, and database theory [9]), electrical engineering (signal processing [SP] [10], estimation and detection, information theory [11], and control theory [12]), operations research [13], and high-performance/scientific computing [14], among others.

In addition to the numerous stages involved in DS, the boundaries between the stages are not clearly defined. The complexity of DS, therefore, calls for a holistic approach, where mathematics, computer science, and engineering, among other disciplines, collectively devise new strategies to address contemporary challenges. This rationale has been shown to be particularly effective when it comes to addressing open problems in DS research, and it therefore should be incorporated in the way DS is taught.

Since the various disciplines that constitute DS research and practice are rarely found together, students, professionals, and even academics have struggled to cobble together pedagogical sources for DS training. As a consequence of this demand, we have more recently witnessed a proliferation of academic/professional programs on DS offered by departments of engineering, business, statistics, or computer science. The significance of DS training and practice [15] has even been identified by policy makers as a key component of national strategies on AI in the United Kingdom, France, Canada, the United States, China, and India, to name a few.

"Data is the new oil"

In Chile, we believe that DS is instrumental to migrate from a natural-resource economy to a knowledge-based one. The Chilean economy, in particular, is largely based on the exploitation of natural resources, such as copper, agriculture, forestry, and fishing. In fact, companies in the mining and agriculture sectors employ a limited (or no) highly trained workforce [16]. Like other developing countries, Chile has been updating its strategy for sustaining productivity and directing its economy toward knowledge-based growth [17]—a process in which data are undoubtedly the raw element [18]. Two popular quotes reflect well the role that DS plays in knowledge-based economies:

- "Data is the new oil. Data is just like crude. It is valuable, but if unrefined, it cannot really be used"—Clive Humby [19].
- "Data science is the sexiest job of the 21st century"—*Harvard Business Review* [20].

These quotes suggest that DS is key for the future, and developing countries, such as Chile, where alternative paths to consolidate the economy are urgently needed, must recognize this opportunity.

The need for DS is, however, not exclusive to the economy but applies to the more general concept of societal development reaching into every realm. The wide-ranging importance of DS has been identified by authorities in Chile, too, where a discussion toward a national strategy on AI is currently underway (see https://www.minciencia.gob.cl/politicaIA), as in many developed societies. However, once the ability to produce data is in place, the next step toward building a knowledge-based society is to equip citizens with tools to extract value from those data. Therefore, teaching DS should be a prime objective for societal development and, thus, a duty of academia.

Scope of the article

The scope of DS's reach and the challenges it presents demand more than a program but, rather, an ecosystem to address its development within the university. This article presents a set of initiatives that encompass teaching DS to engineering students at undergraduate and graduate levels, research interns, project engineers, and professionals. The DS ecosystem described in the following sections is not exclusive to any program or department in our engineering school but, rather, exists as a collection of campus-wide resources that can be extended as demand requires. Our description builds on the teaching and applied research experiences of the authors over the last three years, both independently and collaboratively, with emphasis on blending theory and practice as required in modern, real-world DS challenges.

Our perspective on teaching DS

Academic and professional programs in DS have flourished across the country. This is in line with the global trend, where the success of DS generates a demand for data scientists with educational institutions aiming to fulfill such demand. In a multidisciplinary fashion, available DS programs adopt different perspectives stemming from business, computing, and the natural or social sciences. In this competitive arena, DS programs must establish a clear and focused objective; ours, in particular, is engineering.

We believe that research, teaching, and practice are heavily intertwined in DS: those at the forefront of DS research and practice are the best equipped to teach it. With this idea in mind, our approach adopts an engineering perspective, blending concepts from mathematics and the natural sciences, resources from computer science, and applications of general interest as in [21]. This perspective has been instrumental to equip students with the necessary scientific background while also exposing them to real-world engineering challenges that come from various sectors of human endeavors, including industry, science, and engineering.

Our engineering school in context

FCFM at Universidad de Chile was established in 1842 and (as of 2021) hosts 12 departments and 10 research centers. The civil engineering degree at FCFM, as in most universities in Chile, requires 11 or 12 terms (depending on each specific specialty) of full-time study spanning six years. After this period, our graduates receive both a bachelor of science in engineering (B.Sc.) degree completed after the fourth year and a professional engineer's title (PET) oriented to practical engineering duties. The last credits of the degree focus on elective courses, which can be industry or research oriented, and the thesis work. Critically, some students enroll in master of science (M.Sc.) degree programs in parallel with their last undergraduate year. M.Sc. degree programs at FCFM are two years long, yet the joint B.Sc., PET, and M.Sc. degree can be completed in seven years due to overlapping requirements. In addition to the academic programs, FCFM also offers professional diplomas on different engineering-related topics.

Our DS curriculum builds on the following units at FCFM:

- *The Center for Mathematical Modeling (CMM)*: Areas include probability, optimization, statistical machine learning, machine learning for health care, and scientific computing.
- *The Department of Computer Science (DCS)*: Areas include data mining, natural language processing (NLP), database theory, deep learning (DL), multimedia databases, information retrieval, semantic web, and data compression.
- *The Department of Electrical Engineering (DEE)*: Areas include SP, computational intelligence, robotics, information theory, and control systems.

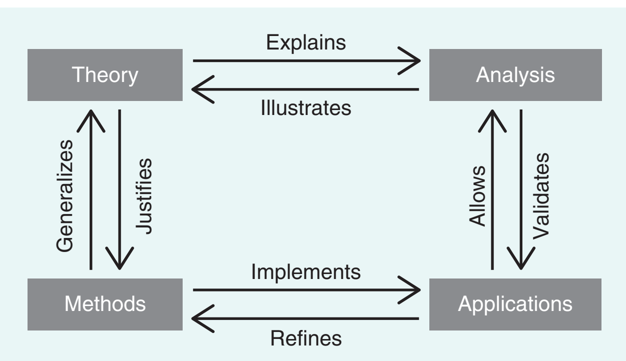

FIGURE 1. The learning objectives and their connections. Theory supports the development of methods and ways of generalizing them to different scenarios. At the same time, applications provide insight into the choice and enhancement of methods, while the results of applications allow for analysis, which validates or refutes the theory on which the other learning objectives are based.

Learning objectives

We integrate the engineering perspective into teaching DS by pursuing four learning objectives. The first, rooted in theory, relates to understanding the data-generating systems to identify challenges and envision solutions at a conceptual level. This can be achieved from first principles or from a data-driven, application-agnostic, machine learning perspective. The second objective relates to selecting the methods for the different stages of DS. This is fundamental for practitioners who should be able to discriminate which tools are appropriate for each task, both at the level of data handling and knowledge extraction.

The third objective relates to applications, through which professionals deal with real-world DS challenges of different natures by formulating the problem in a DS setting where attaining a solution is feasible. The final learning objective focuses on analysis, in which professionals are expected to interpret the results so as to 1) provide explanations, 2) identify possible shortcomings of the methods employed, and 3) explore solutions for such shortcomings based on theory.

Our learning objectives are interconnected, and they support one another, as illustrated in the diagram in Figure 1. Finally, it is worth noting that our objectives are in line with other criteria for accreditation of engineering programs, such as those of the United Kingdom's Accreditation of Higher Education Programs, set by the Engineering Council [22, p. 10] and also by the Institution of Engineering and Technology [23, p. 8].

Components of the DS ecosystem

M.Sc. degree programs

The bases of our DS ecosystem are the M.Sc. programs at the units mentioned in the "Our Engineering School in Context" section: the M.Sc. degree in mathematical modeling, M.Sc. degree in computer science, and M.Sc. degree in electrical engineering, all of which feature a DS specialization and are accredited in the country by the corresponding institutions. Hosted at different departments, these programs offer complementary views, where students are exposed to courses, students, and faculty of different departments. This collaborative environment allows our students to build their own DS profile by mixing resources from different perspectives.

Postgraduate courses

We describe DS-related courses using a two-level categorization, where categories (content and focus) and subcategories are not necessarily mutually exclusive. The content category points to the elements that are taught in each course and follows from our first three learning objectives in the "Learning Objectives" section. The focus category relates to the DS stages toward which each course aims. The complete list of courses can be found in Table 1.

Content

Theory courses are oriented to the formulation and analysis of mathematical and computational models for data management and data analysis. Theory-based courses enable students to understand the limitations of off-the-shelf methods and question existing solutions. Topics for such courses include probability, statistics, stochastic processes, optimization, algorithmic complexity, discrete mathematics, database theory, and dynamical systems.

Methods courses focus on addressing practical DS challenges in a problem-driven fashion. The core of DS methods includes NLP, DL, SP, nonparametric Bayesian inference, transform-based analysis, spectral analysis, Monte Carlo simulation, and data visualization.

Finally, applications courses ensure that students are not only knowledgeable on theory and methods but can also implement them on arbitrary-domain challenges. This is particularly useful for graduates working in a DS-as-a-service environment, such as business analytics, health, climate, or astronomy. To meet the wide range of student interests and needs, we place particular attention on recreating realistic DS scenarios within our courses so that students face all stages of a real-world DS project, i.e., from data management (acquisition, curation, and processing) to data analytics (mining, inference, decision making, and interpretation).

Focus

A course can focus on one or more of the following aspects:

- *Data management*: topics of data handling, such as acquisition, processing, governance, architecture, storage, security, privacy, quality, and curation
- *Data analytics*: knowledge extraction tools, such as machine learning, probability models, statistics, time series, and data mining
- *Application domains*: areas that rely on DS resources, such as text processing, speech synthesis, computer vision, image processing, robotics, econometrics, astrostatistics, and bioinformatics, among others
- *Related fields*: disciplines, not necessarily associated with DS, where knowledge extraction is also relevant (information theory, stochastic simulation, ergodic theory and dynamical systems, and algorithms) or those that provide the skills necessary for data scientists (optimization, algebra, algorithms, and stochastic processes).

Though these courses are offered as a part of the aforementioned M.Sc. degree programs, they are available for all of the students at FCFM, provided they meet course requirements. Additionally, as part of different programs, some course content may overlap (e.g., machine learning and computational intelligence); however, despite this redundancy, the courses have gained considerable popularity of late. Figure 2 shows that the number of students has increased over the last three years for the flagship courses. It is worth mentioning that the contents of these courses have evolved and will continue to evolve over time based on feedback from students and the current state of the art in the field. For some of the courses, the content is publicly available, such as those in Table 2.

Internships

Our DS offerings include internships oriented to professional or research work. Those in professional internships, i.e., interns working on our projects or those of our collaborators, can delve deeper into the practice of DS and, thus, make informed decisions when choosing a DS career.

Research internships provide a unique opportunity for students in the transition undergraduate, M.Sc. or Ph.D. degree programs; students interested in a temporary (usually summer) position; and even visiting students joining mainly from our partner institutions. Research internships provide students with firsthand experience in a DS research laboratory. Funding for research interns (both domestic and international) has been possible through faculty research funds, collaboration networks, and also internationalization grants.

Innovative aspects of our approach

Though it is the practical advantages of DS that usually motivate students to seek training in the field, a necessary step to become proficient in DS is to understand the required theory. As pointed out in [21, Sec. 2.5], however, theory cannot be delivered in a raw manner for DS students as in classical scientific degrees, but, instead, it should be presented in a problem-driven fashion. Additionally, even for those students who are familiar with the theory already, e.g., those holding a degree in mathematics, making the transition from theory to DS practice can be a challenge. Current information technologies, computational resources, and public data sets allow us to offer an ad hoc pedagogical presentation of the theory and its connection to DS practice.

These innovations are key in our ecosystem, just as they have been for teaching statistical SP [24] or AI [25]. We next describe the innovative features in our DS ecosystem (either exploratory or consolidated) and how they enrich the pedagogical process.

Online code repository

As a companion to some of our courses, we include the pedagogical material in a public repository. This allows students to have instant access to lecture notes, slides, assignments, demonstrations, and, in some cases, video lectures. This way of distributing the material has proven advantageous for several reasons. First, both the lecturer and teaching assistants can simultaneously edit the material, minimizing the number of conflicts and maintaining a history of past versions. Second, should last-minute changes occur in the course materials, the up-to-date version is automatically available to the students. Third, the students can visualize the course contents online, which is of particular interest for Jupyter notebooks (JNs)

Table 1. The number of students corresponds to the last term the course was given. The course level can be identified from its code: 4000 or less corresponds to undergraduate courses, 5000 corresponds to both undergraduate and postgraduate classes, and 6000 courses and above are restricted to postgraduate students.

Course	Code	Students/Term	Department	Content			Focus			
				Theory	Methods	Applications	Data Management	Data Analytics	Application Domains	Related Fields
Statistics	MA3402	48	CMM	✓	✓	·	·	✓	·	·
Stochastic Simulation	MA4402	29	CMM	✓	·	·	·	·	·	✓
Machine Learning	MA5204	90	CMM	✓	✓	·	·	✓	·	·
Advanced Machine Learning	MA5309	7	CMM	✓	✓	·	·	✓	·	·
Probability and Statistics for DS	MA5406	12	CMM	✓	✓	·	·	✓	·	·
Lab of Mathematical Modeling	MA5500	7	CMM	·	·	✓	·	·	·	✓
Optimization for DS	MA5705	21	CMM	✓	✓	·	·	·	·	✓
Scientific Computing	MA6201	8	CMM	·	·	✓	✓	✓	·	·
DS Laboratory	MA6202	27	CMM	·	✓	✓	✓	✓	·	·
Algorithms and Data Structures	CC3001	186	DCS	✓	✓	·	·	·	·	✓
Databases	CC3201	117	DCS	✓	✓	·	✓	·	·	·
Introduction to Data Mining	CC5206	70	DCS	·	✓	✓	✓	✓	·	·
Information Visualization	CC5208	17	DCS	·	·	✓	·	✓	·	·
Massive Data Processing	CC5212	60	DCS	·	✓	✓	✓	·	·	·
DS Project	CC5214	17	DCS	·	·	✓	✓	✓	·	·
Image Processing and Analysis	CC5508	14	DCS	·	✓	·	·	✓	✓	·
Pattern Recognition	CC5509	11	DCS	·	✓	·	·	✓	·	·
Business Analytics	CC5615	38	DCS	·	✓	✓	✓	✓	·	·
Deep Learning	CC6204	118	DCS	✓	✓	·	·	✓	·	·
Natural Language Processing	CC6205	57	DCS	·	✓	✓	·	✓	✓	·
Web of Data	CC7220	36	DCS	·	✓	✓	✓	·	·	·
Signal Processing	EL4101	11	DEE	✓	✓	·	·	✓	✓	·
Computational Intelligence	EL4106	51	DEE	✓	✓	·	·	✓	·	✓
Neural Networks and Information Theoretic Learning	EL7006	11	DEE	✓	✓	·	·	✓	·	·
Introduction to Digital Image Processing	EL7007	16	DEE	·	✓	·	·	✓	✓	·
Advanced Image Processing	EL7008	36	DEE	✓	✓	·	·	✓	✓	·
Fault Diagnosis and Failure Prognosis	EL7014	19	DEE	·	✓	·	·	✓	✓	·
Robotics, Sensing, and Autonomous Systems	EL7021	4	DEE	·	✓	·	·	✓	✓	✓
Information Theory: Fundamentals and Apps	EL7024	26	DEE	✓	·	·	·	·	·	✓

(see the next section) that require a specific interpreter in local machines. Fourth, all course materials can be made universally available to the general public beyond our institution. See Table 2 for an example of our courses GitHub repositories.

Interactive programming

Computer programming is best taught with a learning-by-doing approach [26]. For DS in particular, the JN has revolutionized the way we program [27], with a clear impact when teaching and prototyping: it is free, open source, interactive, intuitive, and supported by a strong online community. Our courses feature programming modules on Python (for a majority of courses), R (statistics), MATLAB (electrical engineering), and C++ (scientific computing). Therefore, as JNs are compatible with all of these languages, they are used for in-class demonstrations, in which the lecturer can produce and run examples on the fly.

Additionally, these JNs are distributed to the students for personal study, allowing them to complete, modify, and run examples at their convenience while also exploring creative variants; this is especially required for assignments and project-oriented activities within the courses. Beyond methods-based courses, where the value of code demonstrations is clear, we have learned that JNs and similar software constitute an excellent complement for theory-based courses, too. For example, illustrations of hypothesis testing in statistics and stochastic gradient descent in optimization greatly benefit from modifiable demonstrations when compared to old-school blackboard illustrations.

Evaluations promoting independence and creativity

To a large degree, course evaluations condition the design of the course and its success in transmitting knowledge. A successful evaluation becomes particularly challenging in the context of our teaching objectives (outlined in the "Learning Objectives" section), which aim at having students equally comfortable with both theory and practice. As DS challenges require creative, out-of-the-box solutions, we strive to recreate these requirements in our evaluations.

Whenever the topics allow it, in addition to the theoretical/practical parts, our evaluations incorporate open-ended questions whose objective is to encourage students to build on the concepts learned in the lectures. In these instances, students are required to specify the question and solve it, derive alternative solutions to those problems examined in class, or review the literature for material that has been hinted at in class yet not thoroughly reviewed. In this way, we aim to ensure that students solve realistic problems rather than (just) implementing an off-the-shelf method. This has proven to be particularly challenging for inexperienced students used to well-defined problems that often have a unique solution; these students require close supervision.

Project-oriented learning

In most of our courses, the final evaluation requires students to form groups (of 2–4 members) to complete a project, which

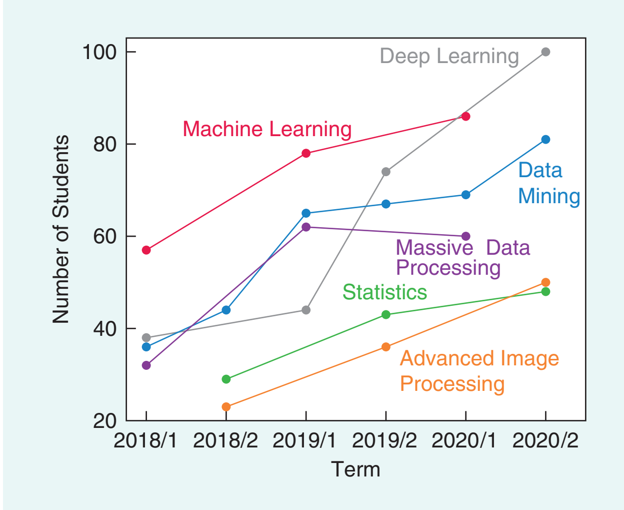

FIGURE 2. The evolution of registered students in the most relevant DS courses over the last three years.

Table 2. The GitHub repositories of some of our courses containing lecture notes, slides, exercises, and demos.

Course	GitHub Address
Machine learning	https://github.com/GAMES-UChile/Curso-Aprendizaje-de-Maquinas
Statistics	https://github.com/GAMES-UChile/Curso-Estadistica
Deep Learning	https://github.com/dccuchile/CC6204
Natural Language Processing	https://github.com/dccuchile/CC6205

can be of theoretical or applied content, or a combination of both. The execution of such a project is developed throughout the course, alongside lectures, where preliminary advances of the projects are monitored as partial course evaluations (tests and assignments). We usually provide students with a repository of project topics built from past courses, industrial projects, and the lecturers' own research portfolio. However, students are also encouraged to propose project themes motivated by their thesis work, entrepreneurial activities, or other topics of interest that can be addressed using DS. We have noticed that it is precisely those projects brought by the students that turn out to be the most successful, most likely because there is a genuine motivation to work on these rather than on a generic assignment.

Communication skills

DS engineers work in interdisciplinary teams and must be able to communicate clearly. Our courses consider four practices aimed at developing these skills. First, the format for the submissions (assignments and reports) is evaluated in terms of presentation, conciseness, clarity, and readability. What students have found particularly challenging here is to constrain their description to a limited number of pages. Second, in seminar-based courses, we use the flipped-classroom method [28], where students teach their fellow classmates. Third, in project-oriented courses, students work on a DS problem for which they must formulate and define the scope, select the methods and strategies to be used, and analyze the results.

In all of these stages, the students work as a team: regular meetings and presentations are conducted to evaluate the ability of the students to communicate the project's state to the rest of the class and the instructor. Fourth, for those courses featuring real-world projects from industry, students present their (finalized) DS projects in a 10-min pitch talk to the company that proposed the challenge, thus validating the communication abilities of the students with actual industrial counterparts.

Early research training

The M.Sc. degree programs described in the "Master of Science Degree Programs" section culminate with a two-semester research thesis that can be of either an applied or theoretical nature; for most students, this experience constitutes their first exposure to research. In their theses, students join other research students and one or more supervisors, sometimes in collaboration with partners from industry, the public sector, or other sciences. Exposing our students to research may also improve their employability: though the majority of our graduates join the industrial sector, their research experience makes them valuable assets in modern industry, which often values research.

Lastly, for those of our graduates who pursue an academic career, the M.Sc. thesis provides a fertile environment for theoretical research, where students join Ph.D. students and postdocs to work under the close supervision of their mentors and, in most cases, successfully publish their findings (see, e.g., [29]–[33]).

Professional and academic development

The public and private (also known as *professional*) sectors as well as academia have witnessed the practical advantages of DS and are eager to understand and incorporate such techniques. We have addressed the demand of these sectors for DS training by transferring our experience from undergraduate and graduate programs at FCFM to the professional domain. We next describe the elements of our DS training focused on professionals.

Professional diplomas

Continuing education courses are the most popular destination for professionals seeking DS training. To stand out from the abundance of offerings from other institutions, the distinguishing feature of our professional diplomas follows from our learning objectives and master's degree programs (see the "Learning Objectives" and "Master of Science Degree Programs" sections) to provide an alternative that blends theory and practice in a problem-oriented manner. In particular, we offer two diplomas relevant to DS through the DCS (see the "Our Engineering School in Context" section): the DS diploma and the more advanced AI diploma. Each of these diploma programs features three evening lectures per week, which are completed over a five-month period.

Both diplomas target professionals from the areas of engineering and science, such as astronomers, geologists, biologists, and engineers, although sociologists and lawyers have also successfully completed the courses. On one hand, the DS diploma focuses on the analysis and handling of complex and massive data sets; the main topics are those related to the fundamentals of databases and data mining, basic statistical tools, big data, information retrieval, and visualization.

The AI diploma, on the other hand, focuses on a more experienced audience (e.g., those graduated from the previous diploma) to train them to 1) lead projects that involve complex and heterogeneous data sources in various forms (e.g., text and images) and 2) effectively communicate and justify their findings. Accordingly, this second diploma features more specific contents, such as DL, evolutionary algorithms, image processing, NLP, and robotics. Finally, both diplomas feature a final project through which students tackle a challenge relevant to their own workplace under the supervision of an academic staff member.

On-demand courses

We have also developed tailored courses for those in the professional sector who currently work in DS. These courses have been offered through CMM (see the "Our Engineering School in Context" section) to partners in banking, mining, and nongovernmental organizations (NGOs) that aim to acquire specific and advanced DS skills. For these courses, the syllabus is jointly designed with the interested party with their particular needs and challenges in mind. The courses work as a blend between a diploma and scientific consultancy, whereby the class demonstrations utilize data provided by the institution. In this way, students learn the impact and shortcomings of standard methods as well as the necessity for developing new tools in a familiar environment and with a clear (problem-driven) purpose.

As a consequence of the interdisciplinarity of DS teams in industry, a recurrent challenge in these on-demand courses comes from the heterogeneous levels of expertise in DS found among students in the same group. This justifies the development of courses for small groups of students with purpose-specific content. In fact, when done face to face, we have found that groups of approximately 15 students in weekly sessions of 2–3 h (with a break) are an appropriate format. This allows us to assess the evolution of the students via discussions in class rather than relying on strict evaluations, which are usually incompatible with the availability of the students in these programs. Additionally, these courses employ most of the innovations described in the "Innovative Aspects of Our Approach" section, especially those related to the demonstrations using JNs and GitHub repositories.

Due to the coronavirus outbreak and the sustained lockdown measures during 2020, we offered our tailored courses in an online format. These have been particularly useful for mining companies, where engineers both in Santiago and close to the extraction sites have taken part in the courses. When implemented remotely, we combined online classes and offline content capsules for students to manage at their convenience.

Training of data engineers

The training of DS engineers involves a wide range of coursework, experiences, and exposure to different sectors and professionals both within academia and industry. The data engineers who go through our programs have the opportunity to work on our research projects alongside principal investigators, research assistants, postgraduate students, and interns; for technology transfer projects, the teams usually comprise engineers, (data) analysts, designers, and software developers. In our collaborative teams, project engineers are constantly exposed to research practice and, depending on the nature of the project, they even participate in applied research publications (see, e.g., [34]–[36]). Furthermore, many of the engineers who take part in our programs will have the opportunity to teach and train others, which further enriches their own development and keeps them current on the latest trends and advances in the field.

Within our laboratories and centers, providing hands-on training in which we tackle problems from different industries not only provides an invaluable DS experience for our project engineers, but it also reinforces our role as educators, bridging the gap between academia and the professional sector. Additionally, when our former engineers move on to find jobs outside academia, our collaboration network is strengthened, and our new project pipeline is also often extended.

Indeed, the considerable demand for DS professionals makes the DS job market quite dynamic: project engineers move from academia to industry and also between companies perhaps more than in other disciplines. One reason for this is that DS projects are often completed on a contract basis in which engineers are hired for a specific time to solve a specific problem. This practice is supported by the perspective that sharing talent across different sectors and job mobility are regarded as positive for career development in the field of DS [15]. This job mobility further underscores the need for our programs to provide a wide exposure to the types of problems and experiences DS engineers are likely to encounter once they leave our centers.

Outreach

In addition to the formal treatment of DS in academia as well as the private and public sectors, we have worked toward making DS advances available to the general public. We consider this to be part of the role of universities in the democratization of knowledge [37], which, in our case, relates to promoting literacy in DS. This can be achieved by scientific dissemination activities organized by the public sector or NGOs. In particular, owing to contemporary online teaching practices, outreach also needs to occur through (virtual) talks, discussion panels, and webinars, which can be backed up on a video (e.g., YouTube) repository site.

Furthermore, to raise awareness about the impact of DS on industry leaders and policy makers, we have held "Data Days," a series of discussion panels organized since 2018 where participants discuss a particular DS topic with an influential invited expert. Topics considered so far have been clinical text mining, digitalization of education, climate change and biodiversity, social organization and representativeness, and the Internet of Things in health care. These panels are designed to encourage discussion between attendees and experts so as to identify opportunities and challenges related to the modernization of the local economy via DS.

Another initiative for disseminating the advances and impact of DS is through seminars addressed to high school and university students. For secondary students in particular, we have seen that DS and AI are becoming popular. In this sense, the Explora outreach program (driven by Chile's Ministry of Education; see https://www.explora.cl/) invites secondary school students to develop a project under the supervision of a DS expert. In fact, some high schools have instituted (Python) programming courses through which our researchers have carried vibrant interactions. Finally, it is relevant to mention that there are initiatives that aim to reduce the gender gap in science, technology, engineering, and mathematics disciplines, and many of these events have recruited DS experts to give open talks or serve as judges in DS competitions.

Open challenges

Teaching DS focuses on shaping highly skilled technical professionals to develop and implement methods to extract information from data in various domains. However, being in close connection to AI, the discipline of DS is also at risk of being automated itself. Therefore, the following question arises naturally: How should we cope with the replacement of data scientists by machines? It is known that Google, Amazon, and IBM provide cost-efficient, modular DS solutions that are the choice of companies relying on DS as a service (DSAAS); it is critical that our graduates can deal with and adapt to the massification of DSAAS. To this end, our graduates should master the underlying theory of DS practices so that they truly are data scientists and not mere DS practitioners and can adapt to changing circumstances in their field of expertise.

Another challenge to be faced by our graduates is that of the so-called social value of the data. Novel tools for data processing have allowed us to identify their value as a means to multiple ends, such as marketing, political campaigns, public policies, and insurance. There are, of course, companies that support their activity purely on the value of data, such as Twitter and Facebook. With the sophistication of DS tools to extract information from data, we are facing an era where—to an extent—data can be considered a commodity; this scenario opens both negative and positive opportunities.

First, how can we guarantee that a small country such as Chile is able to protect its data, when large international conglomerates are at play? For instance, Chile's recently launched Data Observatory (https://www.dataobservatory.net/) will be hosted at Amazon Web Services, which has implications unknown to the public at the time of this writing.

Second, are we able to take the leap forward into a modern technological society by both properly curating our data and developing tools to extract knowledge from them? There are case studies of which our students should be aware in this regard, such as that of Cambridge Analytica [38]. As a consequence, our DS curriculum should feature courses dedicated to the issue of data value and privacy so that our graduates, in addition to being experts on DL, scientific computing, and probabilistic modeling, are also knowledgeable of the value and impact conveyed by the DS tools they handle.

Summary

Developing the described teaching ecosystem has been an enriching experience both as researchers and educators. Through this article, we have highlighted the considerations and innovative aspects we consider meaningful and essential in putting together effective DS curricula for undergraduate and postgraduate students, professionals, and the general public. There is a common denominator in designing a DS ecosystem: finding the appropriate balance among theory, methods, and applications. This interplay is essential to achieve an educational experience for students that is practical (an up-to-date presentation of techniques and solutions), meaningful (covering the advantages and limitations of the strategies and methods), and fundamental (promoting critical thinking and a level of abstraction that facilitate innovation and creativity in DS).

Considering our approach to teaching DS as an ecosystem rather than as a single curriculum has allowed us to widen our scope and include not just courses but resources and outreach initiatives. We hope that the material presented here can help others in the process of developing their own DS programs and ecosystems.

Acknowledgments

We are grateful to FCFM, Universidad de Chile, for its continuous support of DS initiatives. In particular, we deeply thank our colleagues at the DEE, DCS, and CMM, as they have made possible the DS ecosystem described here. This work was supported by the following Agencia Nacional de Investigacion y Desarrollo (Chile) grants: AFB170001 CMM; AFB0008 Advanced Center of Electrical and Electronic Engineering; ICN17_002 Millennium Institute for Foundational Research on Data; and Fondecyt-Iniciación projects 11171165, 11200290, and 11201250.

Authors

Felipe Tobar (ftobar@dim.uchile.cl) received his Ph.D. degree from Imperial College in 2014. He is a researcher with the Center for Mathematical Modeling, Universidad de Chile, Santiago, 8370456, Chile, where he teaches statistics and machine learning courses to undergraduate, graduate, and professional audiences and is the coordinator of Universidad de Chile's Master of Data Science. He also coleads the Data Science Unit at the Center for Mathematical Modeling and has been an associate researcher at the Advanced Center for Electrical and Electronic Engineering, Universidad Santa María, since 2020. His research interests include statistical machine learning and signal processing. He is a Member of IEEE.

Felipe Bravo-Marquez (fbravo@dcc.uchile.cl) received his Ph.D. degree in machine learning from the University of Waikato, New Zealand, where he also held a research fellow position for two years. He is an assistant professor with the Department of Computer Science at the University of Chile, Santiago, 8370456, Chile, and a Young Researcher at the Millennium Institute for Foundational Research on Data. His research interests include the acquisition of knowledge and information from natural language, spanning the overlapping fields of natural language processing, machine learning, artificial intelligence, and information retrieval.

Jocelyn Dunstan (jdunstan@uchile.cl) received her Ph.D. degree in applied mathematics from the University of Cambridge. She holds a joint position with the Engineering and Medicine Faculties of the University of Chile, Santiago, 8370456, Chile, where she has been designing and teaching courses that combine theoretical aspects with real-world applications, especially considering Chilean needs. She leads the Group on Natural Language Processing at the Center for Mathematical Modeling. Her research interests include the application of machine learning and natural language processing to the medical domain.

Joaquín Fontbona (fontbona@dim.uchile.cl) received his Ph.D. degree in mathematics from Université Pierre et Marie Curie in 2004. He is an associate professor in the Department of Mathematical Engineering and an associate researcher at the Center of Mathematical Modeling (CMM), both at the University of Chile, Santiago, 8370456, Chile. He has designed and taught courses on stochastic simulation and algorithms and leads the CMM Mining Group, where he develops collaborative data science projects between the university and industry. His research interests include probability theory and stochastic modeling, with applications in physics, biology, engineering sciences, and data analysis.

Alejandro Maass (amaass@dim.uchile.cl) received his Ph.D. degree in pure mathematics from Université Aix-Marseille. He is a professor with the Department of Mathematical Engineering, and, in 2000, he joined the Center of Mathematical Modeling (CMM), both at Universidad de Chile, Santiago, 8370456, Chile. In 2002 he created the Laboratory of Bioinformatics and Mathematics of the Genome at CMM to provide computational and mathematical tools to solve real-life biotechnological challenges. His research

interests include ergodic theory, focused on symbolic and topological dynamics, and new scientific adventures in mathematical biology and bioinformatics.

Daniel Remenik (dremenik@dim.uchile.cl) received his Ph.D. degree in applied mathematics from Cornell University in 2009. He is an associate professor with the Department of Mathematical Engineering and the Center for Mathematical Modeling (CMM) at Universidad de Chile, Santiago, 8370456, Chile. Since 2014, he has led projects at CMM devoted to the application of machine learning techniques to solve problems in industry and, since 2018, has codirected the Data Science Group at CMM. His research interests include probability theory, with focuses on including interacting particle systems and the problems lying at the intersection with mathematical physics.

Jorge F. Silva (josilva@ing.uchile.cl) received his Ph.D. degree in electrical engineering from the University of Southern California in 2008. He is an associate professor in the Department of Electrical Engineering, University of Chile, Santiago, 8370456, Chile, and principal investigator of the Advanced Center for Electrical and Electronic Engineering, Basal Project FB0008. He was an associate editor of *IEEE Transactions on Signal Processing* (2006–2008). His research interests include information theory, statistical learning, statistical signal processing, learning and decision, compressed sensing, and wavelet and multiresolution analysis. He is a Senior Member of IEEE.

References

[1] T. S. Kuhn, *The Structure of Scientific Revolutions*. Chicago: Univ. of Chicago Press, 2012.

[2] M. Gibbons, Ed., *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Societies*, vol. 1, no. 2, London: SAGE Publications, 2010.

[3] J. T. Klein, *Interdisciplinarity: History, Theory, and Practice*. Detroit, MI: Wayne State Univ. Press, 1990.

[4] D. Donoho, “50 years of data science,” *J. Comput. Graph. Stat.*, vol. 26, no. 4, pp. 745–766, 2017. doi: 10.1080/10618600.2017.1384734.

[5] J. W. Tukey, *The Future of Data Analysis*. New York: Springer-Verlag, 1992, pp. 408–452.

[6] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA: MIT Press, 2012.

[7] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[8] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis*. Boca Raton, FL: Chapman and Hall/CRC Press, 2013.

[9] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases*, vol. 8, Reading, MA: Addison-Wesley Reading, 1995.

[10] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-time Signal Processing*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 1999.

[11] D. J. MacKay, *Information Theory, Inference and Learning Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2003.

[12] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, vol. 1, no. 2, Belmont, MA: Athena Scientific, 1995.

[13] W. L. Winston and J. B. Goldberg, *Operations Research: Applications and Algorithms*, vol. 3, Belmont, CA: Thomson/Brooks/Cole, 2004.

[14] T. E. Oliphant, “Python for scientific computing,” *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 10–20, 2007. doi: 10.1109/MCSE.2007.58.

[15] “Dynamics of data science skills: How can all sectors benefit from data science talent?” The Royal Society, London, May 9, 2019. [Online]. Available: https://royalsociety.org/topics-policy/projects/dynamics-of-data-science

[16] R. Doner and B. R. Schneider, “Technical education in the middle income trap: Building coalitions for skill formation,” *J. Dev. Stud.*, vol. 56, no. 4, pp. 680–697, 2020. doi: 10.1080/00220388.2019.1595597.

[17] OECD and United Nations, *Production Transformation Policy Review of Chile*. Paris: OECD Publishing, 2018. [Online]. Available: https://www.oecd-ilibrary.org/content/publication/9789264288379-en

[18] C. Gutiérrez, “Data, science and society,” *Notes for a talk at LEARN Final Conference, Economic Commission for Latin America and the Caribbean (ECLAC), Senate House*, Univ. of London, London, Tech. Rep., 2017. [Online]. Available: https://arxiv.org/abs/1707.07029

[19] C. Humby, “ANA senior marketer’s summit,” Kellogg School, New York, Tech. Rep., 2006. [Online]. Available: https://ana.blogs.com/maestros/2006/11/data_is_the_new.html

[20] T. H. Davenport and D. Patil, “Data scientist,” *Harv. Bus. Rev.*, vol. 90, no. 10, pp. 70–76, 2012.

[21] R. D. De Veaux et al. “Curriculum guidelines for undergraduate programs in data science,” *Annu. Rev. Statist. Its Appl.*, vol. 4, no. 1, pp. 15–30, 2017. doi: 10.1146/annurev-statistics-060116-053930.

[22] Engineering Council, *The Accreditation of Higher Education Programmes*, 3rd ed., 2014. [Online]. Available: https://www.engc.org.uk/

[23] “Guidance on how to meet the learning outcome requirements for accreditation,” The Institution of Engineering and Technology, London, 2015. [Online]. Available: https://www.theiet.org/media/6274/guidance-on-meeting-ahep-learning-outcomes.pdf

[24] S. Kanna, W. von Rosenberg, V. Goverdovsky, A. G. Constantinides, and D. P. Mandic, “Bringing wearable sensors into the classroom: A participatory approach [SP education],” *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 110–130, 2018. doi: 10.1109/MSP.2018.2806418.

[25] D. Kumar and L. Meeden, “A robot laboratory for teaching artificial intelligence,” *SIGCSE Bull.*, vol. 30, no. 1, pp. 341–344, 1998. doi: 10.1145/274790.274326.

[26] H. W. Reese, “The learning-by-doing principle,” *Behav. Dev. Bull.*, vol. 17, no. 1, p. 1, 2011. doi: 10.1037/h0100597.

[27] J. M. Perkel, “Why Jupyter is data scientists’ computational notebook of choice,” *Nature*, vol. 563, no. 7729, pp. 145–147, 2018. doi: 10.1038/d41586-018-07196-1.

[28] J. L. Bishop and M. A. Verleger, “The flipped classroom: A survey of the research,” in *ASEE Nat. Conf. Proc.*, Atlanta, GA, 2013, vol. 30, no. 9, pp. 1–18. [Online]. Available: https://peer.asee.org/the-flipped-classroom-a-survey-of-the-research

[29] F. Tobar, L. Araya-Hernández, P. Huijse, and P. M. Djuric, “Bayesian reconstruction of Fourier pairs,” *IEEE Trans. Signal Process.*, vol. 69, pp. 73–87, 2021. doi: 10.1109/TSP.2020.3038135.

[30] G. Parra and F. Tobar, “Spectral mixture kernels for multi-output Gaussian processes,” in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 30, 2017, pp. 6681–6690. [Online]. Available: https://papers.nips.cc/paper/2017/hash/333cb763facc6ce398ff83845f224d62-Abstract.html

[31] R. Claveria, R. Mendez, J. F. Silva, and M. Orchard, “Visual binary stars with partially missing data: Introducing multiple imputation in astrometric analysis,” *Publications Astronom. Soc. Pacific*, vol. 131, no. 1002, pp. 1–19, 2019. doi: 10.1088/1538-3873/ab22e2.

[32] P. Barceló, M. Monet, J. Pérez, and B. Subercaseaux, “Model interpretability through the lens of computational complexity,” in *Proc. Adv. Neural Inform. Process. Syst.*, vol. 33, 2020.

[33] J. Dunstan, M. Aguirre, M. Bastías, C. Nau, T. A. Glass, and F. Tobar, “Predicting nationwide obesity from food sales using machine learning,” *Health Informat. J.*, vol. 26, no. 1, pp. 652–663, 2020. doi: 10.1177/1460458219845959.

[34] T. de Wolff, A. Cuevas, and F. Tobar, “MOGPTK: The multi-output Gaussian process toolkit,” *Neurocomputing*, vol. 424, no. 1, pp. 49–53, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0925231220317410 doi: 10.1016/j.neucom.2020.09.085.

[35] C. Delplancke, J. Fontbona, and J. Prado, “A scalable online algorithm for passive seismic tomography in underground mines,” *Geophysics*, vol. 85, no. 4, pp. WA201–WA211, 2020. doi: 10.1190/geo2019-0440.1.

[36] P. Báez, F. Villena, M. Rojas, M. Durán, and J. Dunstan, “The Chilean waiting list corpus: A new resource for clinical named entity recognition in Spanish,” in *Proc. 3rd Clin. Natural Language Process. Workshop*, 2020, pp. 291–300. doi: 10.18653/v1/2020.clinicalnlp-1.32.

[37] G. Biesta, “Towards the knowledge democracy? Knowledge production and the civic role of the university,” *Stud. Philos. Educ.*, vol. 26, no. 5, pp. 467–479, 2007. doi: 10.1007/s11217-007-9056-0.

[38] J. Isaak and M. J. Hanna, “User data privacy: Facebook, Cambridge Analytica, and privacy protection,” *Computer*, vol. 51, no. 8, pp. 56–59, 2018. doi: 10.1109/MC.2018.3191268.

SP

Aleš Procházka, Oldřich Vyšata, and Vladimír Mařík

Integrating the Role of Computational Intelligence and Digital Signal Processing in Education

Emerging technologies and mathematical tools

©SHUTTERSTOCK.COM/LUCKYSTEP

Rapid progress in the development of technological and computational tools has motivated substantial changes in the educational approach to the different disciplines of signal, image, and video processing. Moreover, the parallel evolution of sensor systems, data acquisition methods, and computational intelligence has emphasized the importance of signal processing and information engineering, particularly its role in integrating different scientific disciplines through the use of a common set of tools and underlying mathematics. Modern educational courses follow these trends and generally combine the teaching of fundamental computational methods of signal and system modeling with applications to selected case studies. The unifying idea is to apply similar mathematical methods for data processing in completely diverse areas. Emerging methods used in education contribute to this progress, and they provide opportunities to bring together specialists from different disciplines. New technologies facilitate real or virtual activities through excursions to remote laboratories, allowing the demonstration of robotic and speech recognition systems, for example. Participation in seminars, videoconferences, and discussions during colloquia meetings, when included in educational courses, can form further progressive and attractive teaching methods for the rapidly developing interdisciplinary area of signal processing.

Introduction

Digital signal processing (DSP) has developed into a general interdisciplinary area with a wide range of applications related to the analysis of multichannel and multidimensional signals. These signals may represent any physical, engineering, biomedical, or acoustic variables [1], [2]. Even though applications may cover completely disparate areas, their mathematical backgrounds are generally very similar. In many cases, this allows the processing of vectors, matrices, and multidimensional arrays (representing discrete substitutes of observed signals) with the use of similar computational methods.

DSP methods form a unifying platform for many diverse branches of research. The position of DSP has similarities to the observations made in historical discussions on the differentiation and

Digital Object Identifier 10.1109/MSP.2021.3058634
Date of current version: 28 April 2021

integration of sciences [3], which formed the idea that a too-narrow specialization of scientists makes it harder for them to understand each other. This observation is still very relevant, and it seems that information engineering and signal processing can contribute to the integration of different multidisciplinary topics. In this context, a well-structured educational process has great potential to incorporate different scientific areas and to assist in their understanding.

Fundamental courses on signal processing incorporate sensor systems and their physical principles, numerical analysis, and DSP and image processing methods. Associated tools are sometimes presented in their historical context, which is informative for both students and academics interested in DSP methods. This approach points to the fact that some theoretical inventions predated the development of modern computational systems by many centuries but also contributed to the separate evolution of some research areas, leading to differentiation into new branches of science.

DSP courses have evolved rapidly in parallel with the fast development in computer technologies for human–machine interaction, robotic systems, and assistive technologies, which use different biosensors, data fusion, and wireless communication systems. These developments have also contributed to new innovations in educational tools, which attempt to integrate the view of various signal processing algorithms [4], [5], methods of their presentation, and implementation. There exists a very close interaction between practical needs and theoretical tools, including methods of signal analysis [6], statistical signal processing, digital filtering [7], machine learning, programming methods, and visualization technologies. This multidisciplinary interaction again points to the general issue of effective communication among scientists from different research areas.

To demonstrate the main applications of DSP studies and their evolving role in integrating multiple scientific disciplines, selected case projects with clearly specified applications are being introduced into the classical course concept. This idea is motivated by the very fast recent development of computer technologies and sensor systems [8], which has contributed to a rapid progress in data processing [9], [10].

Standard and online seminars for students from various universities with different backgrounds and interests in information engineering should therefore combine the theory and applications of DSP methods in selected areas. Figure 1 shows a graphical depiction of some courses presented at the University of Chemistry and Technology and the Czech Technical University in Prague, which are used on a website to promote the courses [11]. The appeal of this approach to students was verified at several summer schools and during virtual teaching through videoconferencing systems used for video colloquia. The success of this method is also evidenced in the response to online courses and virtual laboratories [12]–[15] through classroom activity detection monitored by different universities and discussed at virtual conferences. The evidence indicates that this approach is highly appreciated by both students and academics.

The main objective of this article is to present key learning aims and recent experience in the delivery of online courses, computational methods, and their applications. In this context, our goal includes surveying selected case studies that illustrate the interdisciplinary role of signal processing with the use of similar methods in different applications. We also summarize some aspects of modern teaching methods used in signal processing.

Educational methodology

The study of DSP methods is founded on three main pillars: 1) educational tools with face-to-face or online courses, laboratories, and possibilities of data acquisition; 2) computational methods, programming, and visualization tools forming the theoretical background of signal analysis; and 3) applications.

Educational tools

Educational tools for studying signal processing include, at most universities, the use of MATLAB. This software package is used for programming, the use of numerical methods, and the statistical analysis of observed signals.

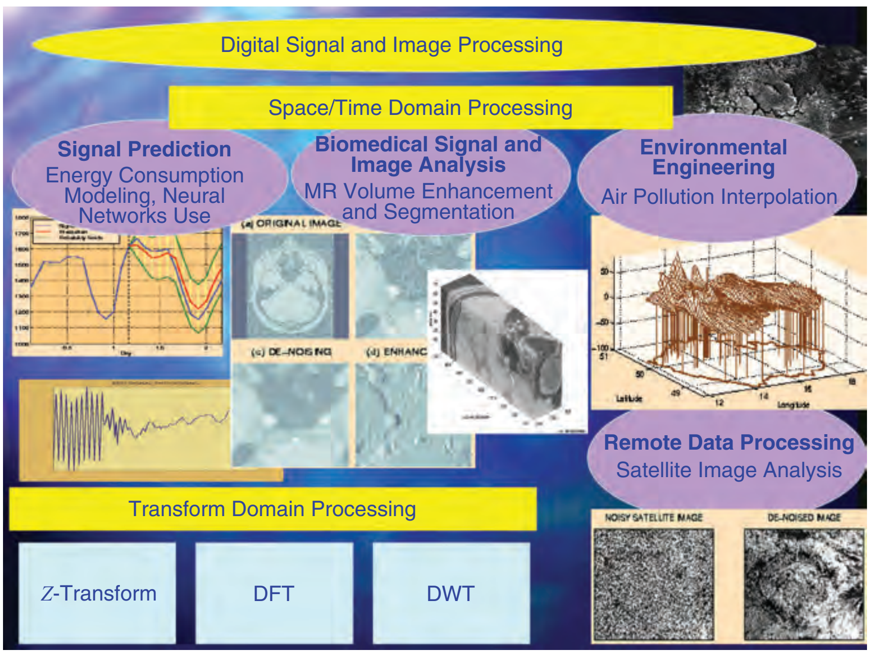

FIGURE 1. The front webpage of a selected intensive DSP course [11] to attract the attention of students to mathematical DSP methods in the space/time and transform domains [based on the z transform, the discrete Fourier transform (DFT), and the discrete wavelet transform (DWT)] and chosen applications related to previous projects and dissertations of our students and devoted to signal prediction, denoising, biomedical signal analysis, and air pollution analysis using terrestrial and satellite measurements.

The direct acquisition of real data from different sensors and camera systems connected to a computer is essential for both students and academics. Processing their own data sets using, e.g., the sensors of their own smartphones increases the students' responsibility and engagement with reliable data acquisition.

The classical format of lectures and video courses is now supported by self-study using ebooks [4], [12] and Internet sources. Completely new virtual tools introduced recently enable the use of videoconferencing methods (including Microsoft Teams, Zoom, Google Meet, Cisco Webex, and Jitsi Meet) that seem to be very convenient for both students and academics. The university license of MATLAB enables students to use all of the toolboxes at home, and a virtual private network allows access to university sources. Students and academics can simultaneously share the same files, discuss proposed methods, and use a computer "whiteboard" to write and draw remarks. Real-time communication can thus be very personal and useful.

These emerging videoeducational technologies can bring new possibilities to the whole educational process. It can be attractive to combine virtual lectures with virtual conferencing (proposed by IEEE for the future) and with video colloquia as well. This approach was introduced very successfully during the ICASSP 2020 conference in Barcelona, and hybrid meetings are planned for the future.

There are several advantages to the use of videoconferencing during the educational process. They include online meetings with experienced lecturers from different countries, who share their presentations during live "webinars." A very popular "night of science" organized online in many cities now enables students to visit real scientific laboratories and motivates their studies of theoretical topics related to different research areas, including signal processing. Moreover, online defenses of dissertations enable the easy participation of reviewers and committee members from different institutes and the establishment of links for further collaboration. But a balance between online and face-to-face education is necessary. There are still formal and legislative restrictions on these activities at some universities, and contact among humans cannot be replaced by robotic systems only.

Computational methods

The interconnection between mathematics and computational tools forms a very attractive background for students of signal processing. The original idea that many numerical objects are associated with multidimensional arrays [16] created an integration platform for MATLAB and other programming languages. Some alternatives now to MATLAB include Octave and Python.

Information engineering and signal processing form a joint platform for several research areas that are making use of data processing techniques. The requirements of real multidimensional signal analysis problems in different applications have motivated profound studies of associated theoretical topics [17], [18]. System modeling and descriptions of signals and systems in the time-frequency and time-scale domains, optimization, and machine learning have formed separate educational subjects. Mathematical topics include different numerical and statistical methods, the discrete Fourier transform (DFT) and discrete wavelet transform (DWT), and methods related to computational intelligence. Special attention has been paid to machine learning methods, feature extraction tools, and different classification algorithms based upon neural networks and Bayesian methods.

Even though the applications differ, it is necessary to analyze vectors or matrices of the observed variables, to apply digital filters for the rejection of undesirable signal or image components, and to use methods for pattern recognition in many cases. Students of signal processing should understand the unifying role of information engineering, which allows a better understanding of scientists in diverse areas of study using similar mathematical tools but different terminologies in some cases. The importance of these methods can also be illustrated during student excursions to different laboratories, including those of robotic and speech recognition systems, to see the purpose of their studies.

The combination of computational methods with visualization tools, 3D modeling, and augmented reality associated with video presentations forms a new rapidly developing area. These applications are very pertinent for studies of specific mathematical methods.

Applications

Case studies have a very important role in studying DSP methods since they point to the integration of this interdisciplinary area based on the processing of data with different applications using similar mathematical tools for their analysis. The most important fundamental data processing methods are mentioned for each selected case study that follows. Case studies enrich DSP courses and provide motivation for learning about the associated theoretical areas.

The interconnection between the practical implementations of signal processing methods and different types of interdisciplinary courses can contribute to very attractive educational activities. Real data and associated student projects support the active participation of students, who can thus see the real implementation of different mathematical methods.

The following selection of projects supervised by authors includes those with different biomedical and neurological applications to motivate studying the associated general DSP methods.

The case studies

The selected case studies include processing of signals acquired by simple mobile sensors (accelerometers, gyrometers), camera systems (red-green-blue, thermal, and depth cameras), scanners, and professional systems [2] used for simple or complex data processing. These studies are related in most cases to student projects and dissertations that were later published, to motivate students to apply their theoretical knowledge of DSP methods to their own research.

The unifying idea behind all case studies is that similar mathematical methods can be used in completely diverse areas. The process of digital filtering for signal or image denoising associated with spectral analysis forms a common basis for many applications. Feature extraction, recognition, and classification are included in many biomedical and engineering problems.

The following selection was used during current DSP courses and is related to biomedical signal processing, motion monitoring, and evaluation of sports activities. But the choice of applications is still modified according to the study program, the results of the students' projects, and the progress of sensor technology each academic year.

Neuroscience and the analysis of brain activities

The analysis of multichannel electroencephalogram (EEG) signals [Figure 2(a) and (b)] recorded by electrodes on the head surface forms the fundamental source of information about brain activities and their disorders. These signals are also used for age-related changes of brain actions [19], as presented in Figure 2(c).

Specific applications include analysis of cognitive functions and mental abilities [20] (see Figure 3), changes in intellectual performance, and human–machine interaction. These topics are closely related to robotic systems, assistive technologies, and computational intelligence.

The methodologies of EEG signal processing include data analysis by the DFT, signal denoising using finite-impulse response filters in most cases, signal segmentation by selected statistical and change-point detection methods, feature detection, and classification.

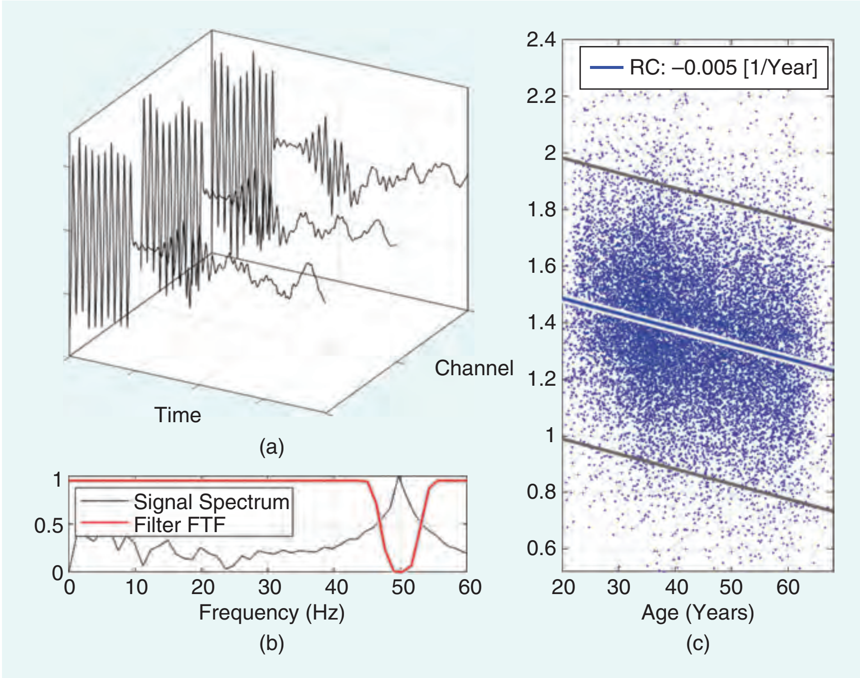

FIGURE 2. An EEG signal analysis presenting (a) the process of signal denoising applied to selected EEG channels recorded with a sampling frequency of 200 Hz, (b) the signal spectrum with its interfering dominant component of net frequency 50 Hz and the FTF of the bandpass filter for its rejection, and (c) results of age-related changes of the λ coefficient in an EEG (as a measure of the interconnection of brain neurons) for a set of 17,722 individuals of different ages, with the associated regression line and its 95% confidence bounds. FTF: frequency transfer function; RC: regression coefficient.

Health science and motion monitoring

Motion analysis is an important topic in the monitoring of physical activities [21], rehabilitation, and recognition of neurological disorders. Current studies include motion assessment using accelerometers [22] inside mobile sensors located at selected body positions and the recording of changes in the heart rate during cycling, under different body loads. Figure 4 presents selected results of cycling monitoring using an accelerometer located at the spine.

The methodologies of motion analysis span a very wide area: they cover the study of different sensors (accelerometers, gyrometers), camera systems, positioning sensors (GPS), and sensors recording associated biomedical signals (such as the EEG and electrocardiogram). Multichannel signal processing then involves the time synchronization of the data sets, their digital filtering, resampling, feature extraction in the transform domains (DFT, DWT), and classification of motion patterns.

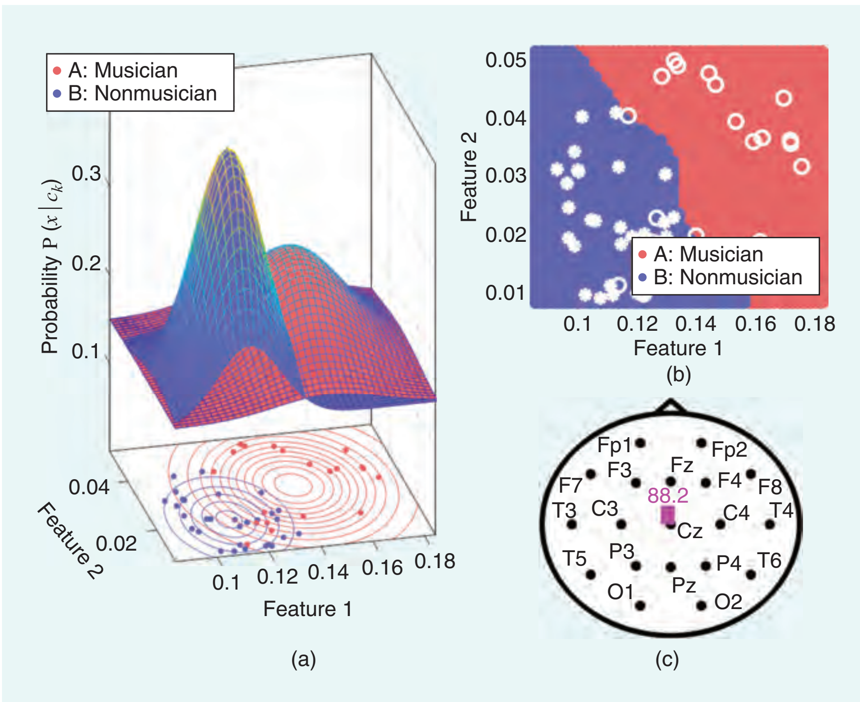

FIGURE 3. An analysis of EEG signals recorded with sampling frequency 200 Hz and its application in cognitive science, presenting (a) the distribution probabilities of foreign language perception for the set of individuals with and without a musical background for the selected electrode, (b) a neural network classification based on data recorded at the selected electrode and chosen signal features, and (c) the location of the most significant electrode with the highest classification accuracy and its value.

Machine–man interaction and gait analysis

The diagnosis of movement disorders, including the detection of gait features [23], [24], forms a very important neurological area closely related to machine–man interaction. It uses images and data from different biosensors, accelerometers, and camera systems. An example of the use of Microsoft Kinect (or a similar device) for the acquisition of gait features and skeletal model construction is presented in Figure 5.

The image and depth sensors of this system enable one to obtain image frames with a given frame rate and detect joints in the 3D space. Inverse kinematics [25] and spatial modeling, employing video systems and the Microsoft Kinect device, can be used in neurology for diagnostic purposes related to the detection of motion disorders, including Parkinson's disease, for rehabilitation, and for studies of human–machine interaction and computational intelligence in engineering.

The methodologies related to DSP include the application of digital filtering, image segmentation and its component recognition, time synchronization of signals recorded by different sensors, feature extraction, and classification. The application of machine learning includes the study of optimization methods, neural networks, and deep learning methods as well.

Thermography and breathing analysis

Current studies include pattern recognition and the analysis of physiological data acquired by thermal camera sensors during rehabilitation. Information related to the distribution of thermal regions (Figure 6) and their evolution over time can then be correlated with the heart rate [26], [27] during different physical activities (on a home exercise bike) and compared with results from the depth cameras.

The methodology of infrared thermal mapping includes the application of image filtering methods, segmentation of image components, and estimation of the evolution over time of selected spectral components. These problems also motivate studies of adaptive methods for the detection of time-dependent regions of interest.

Augmented reality in stomatology

Intraoral scanning technology has brought a completely new approach to dental examinations [28], [29]. In comparison to traditional plaster casts, it allows a more precise digital analysis of dental arch components (see Figure 7) during the treatment of dental disorders. The data acquired can also be used for the creation of 3D models using 3D printers.

The application of signal processing methods in stomatology is very broad, covering the areas of image enhancement, digital filtering, detection of specific image regions, their registration (during the treatment), and methods of 3D modeling. Diffuse reflectance spectroscopy also requires the use of machine learning methods for the detection and classification of dental caries.

Polysomnography and the analysis of sleep disorders

The study of multichannel biomedical signals acquired in a sleep laboratory [30], [31] forms a very interesting research area. An example is presented in Figure 8. The data sets represent mostly polysomnographic overnight records of healthy individuals and individuals with sleep disorders, including sleep apnea and restless legs syndrome. The number of waking and rapid eye movement (REM) stages can be analyzed with respect to sleep disorders and age.

The methodologies used in these data analyses include the application of methods for multichannel signal processing, digital filtering, segmentation, feature extraction, and classification. Classification methods include classical algorithms (decision

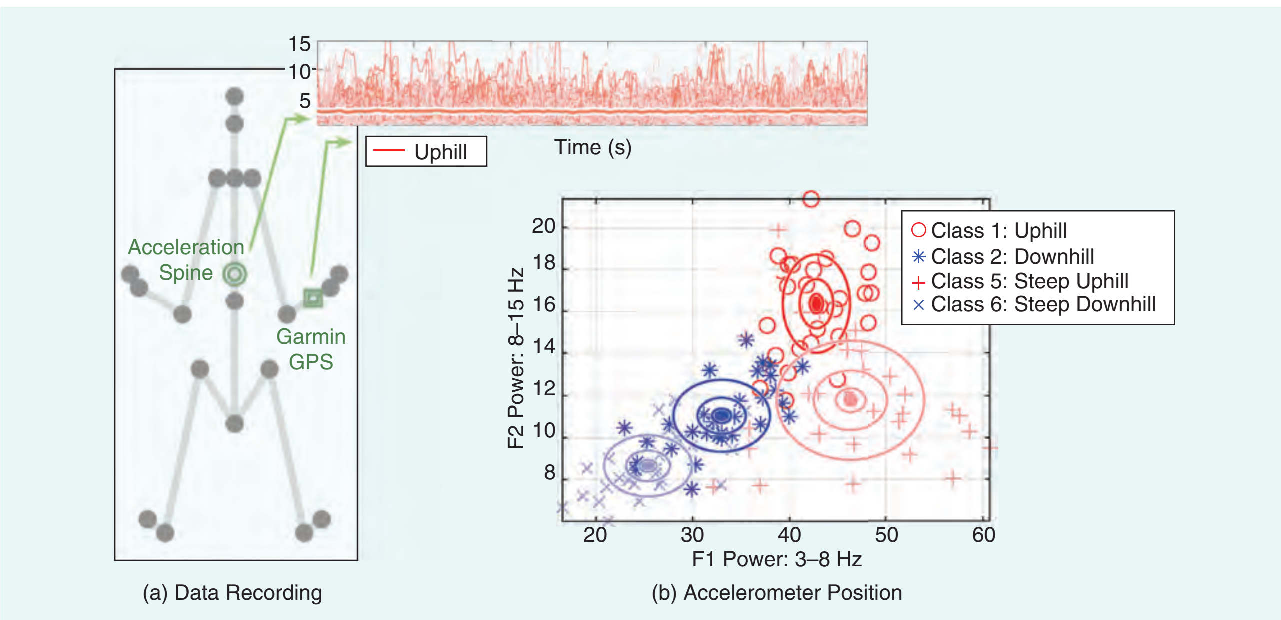

FIGURE 4. An application of DSP in motion analysis, presenting (a) the body positioning of accelerometers and the Garmin system for data acquisition during cycling with a sample accelerometric signal for the sensor located at the spine and (b) distribution of the mean power in selected frequency ranges for different cycling slopes with the centers of gravity for each class and associated regions with selected multiples of standard deviations.

tree, nearest neighbor method, support vector machine, Bayesian methods) and adaptive methods (neural networks, deep learning).

Discussion

Educational courses devoted to computational intelligence combined with specific case studies and physical or virtual excursions to robotic systems and speech recognition laboratories form an attractive way to study signal and image processing methods. Data sharing and their synchronization for analysis, either online or by mobile devices, form another appealing platform for studying numerical methods, DSP, and machine learning [32]. Modern algorithmic tools, including live editors, and the practical implementation of selected methods motivate profound studies of the theoretical background of DSP methods as well.

The possibility of acquiring one's own data sets using modern sensor systems forms another intriguing way to further study DSP. Experiments with physical signals recorded by smartphones motivate the self-study of associated signal analysis tools. This approach forms an alternative to processing professionally recorded signals and images in engineering laboratories, by satellite systems, or in a clinical environment.

The emerging teaching technologies should incorporate the flexibility of DSP courses with the standard structure of their parts devoted to mathematical principles (including studies of signal and system modeling, the DFT, spectral analysis, digital filters, and optimization methods) and case studies that can be different each academic year. This approach, based upon data acquisition with various sensor systems, can have appeal for both academics and enthusiastic students. The rapid pace and variety of technological progress reduces the danger of any stagnation in the development of the content of these courses.

The proposed structure of the educational courses was verified by experience from many courses presented, over more than 30 years, to students with diverse backgrounds, including international students from different countries. The number of students in each course (devoted to mathematical methods, DSP and image processing, and computational intelligence) is usually limited to 20 for better personal contact.

Table 1 lists a few principal subjects presented over the last 30 years for undergraduate and research students from the University of Chemistry and Technology and the Czech Technical University in Prague. These courses included in the complete curricula [11] are devoted to MATLAB programming, numerical and symbolic methods, signal and image processing, optimization, machine learning, and multimedia signal processing [33].

The explanation of selected methods presented during DSP courses [11] has been combined with specific research projects devoted to simulated and real signal processing. The formal structure of these courses is similar and includes

- interactive lectures in computer laboratories in which the explanation of new topics is associated with verification of algorithms by students on their own computers
- presentation of specific case studies introduced in some instances by research students to motivate undergraduate students for their own future scientific research
- computational laboratories with studies of specific methods and possibilities of students' acquisition of their own data sets
- development of several individual student projects each term, with their recommended structure close to that of scientific

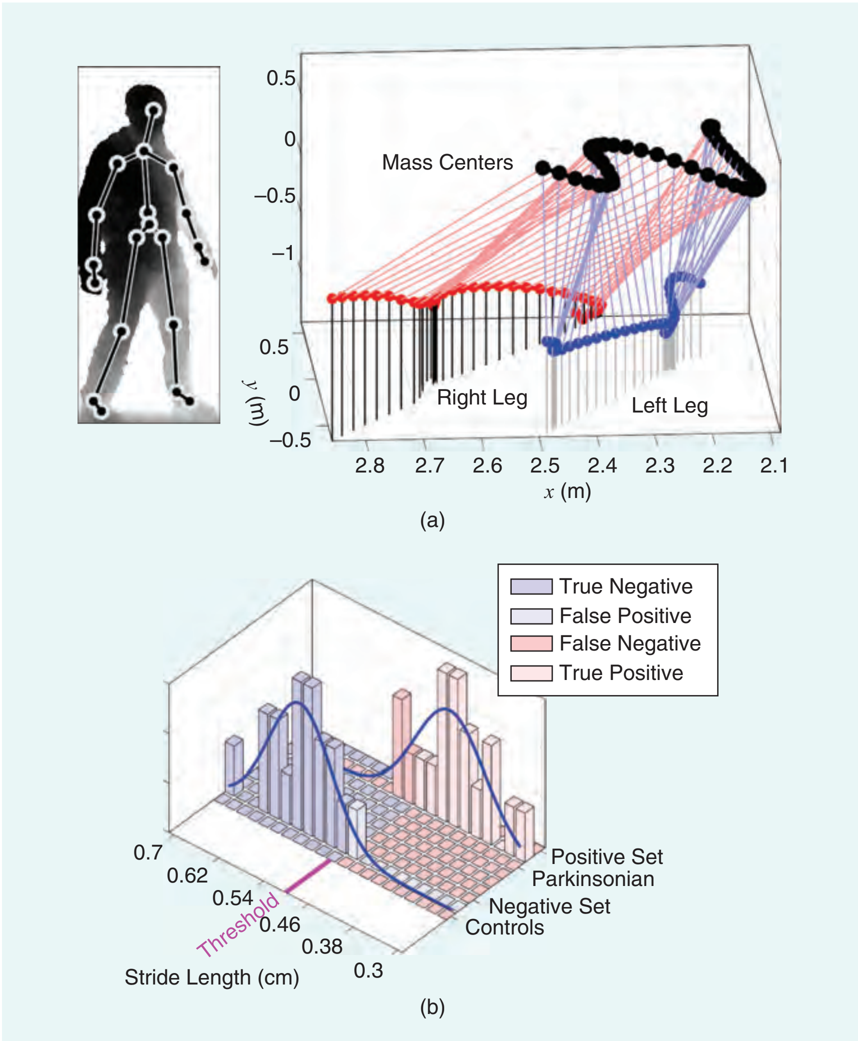

FIGURE 5. An analysis of data sets from a depth camera acquired with a given frame rate and camera resolution, with the application to spatial modeling, presenting (a) the depth matrix with a skeletal model from the Microsoft Kinect system and the evolution of mass and leg centers for a selected walk segment and (b) the histogram of the stride length for individuals with Parkinson's disease (positive set) and age-matched controls (negative set) with the distributions of true negative/positive and false negative/positive values for a selected criterion value.

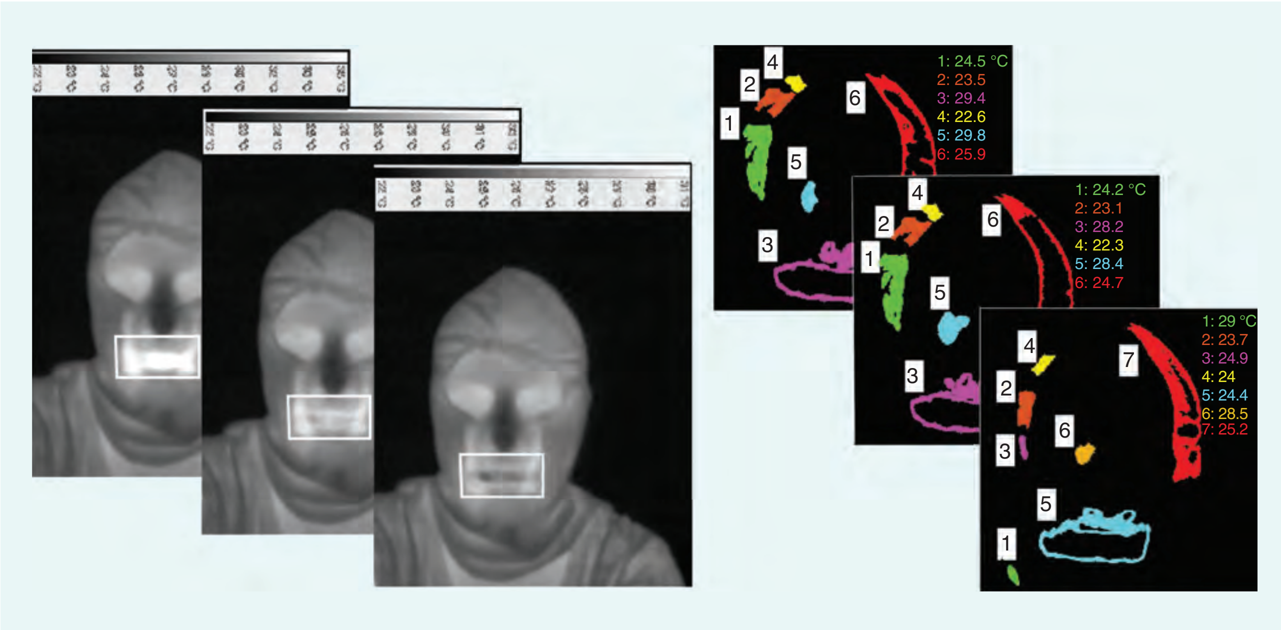

FIGURE 6. Adaptive image processing applied to segmentation of video sequences recorded by a thermal camera to detect breathing regions in each video frame of the facial area, presenting the set of the subsequent thermal images with associated thermo bars and contour plots of changing isothermal regions.

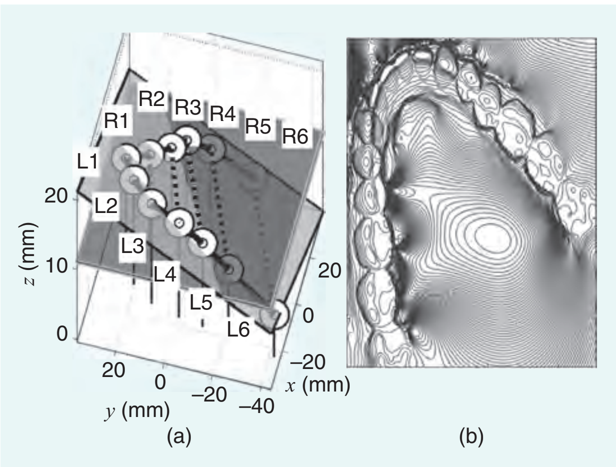

FIGURE 7. A spatial analysis of the lower dental arch used to measure positions of separate objects after an operation, presenting (a) the spatial location of teeth centers during dental-plane rotation into the horizontal position and (b) the corresponding contour plot used for its evaluation, 3D modeling, and observation of the healing process after surgery. R: right; L: left.

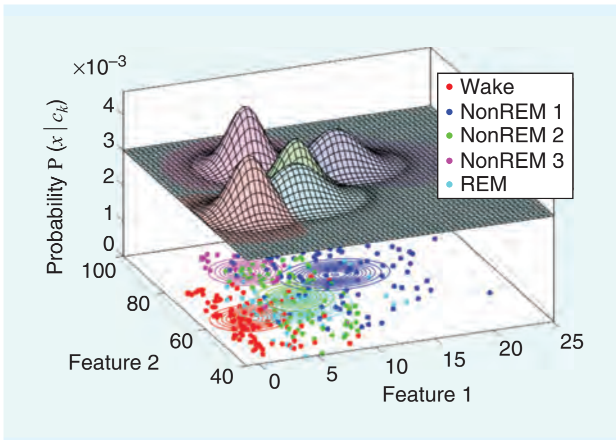

FIGURE 8. The use of EEG signals recorded during the night with a sampling frequency of 200 Hz for analysis of sleep stages presenting the distribution probabilities associated with five classes (Wake, NonREM1, NonREM2, NonREM3, REM) for features evaluated as the mean power in two frequency bands (Feature 1/EEG: 4–8 Hz and Feature 2/breathing: 0.15–0.25 Hz) of a selected overnight record, chosen electrode, and segments 30 s long. REM: rapid eye movement.

papers and using professional tools (like MATLAB and LaTeX), to prepare students for their own later research and publication activities

- final exams or a colloquium.

This structure was enabled by university licenses of the educational software, which allowed students to work at home. Document sharing (using, for instance, the Overleaf system) forms another attractive educational tool. The structure of these courses is similar both for standard and video meetings in the selected virtual environment. In all cases, a classical whiteboard or computer

Table 1. Selected subjects related to DSP with the number of teaching hours per week and students enrolled during academic years 2017–2020.

Subject Name	Extent [h/Week]	Number of Students (2017–2020)
Mathematical Methods in Engineering	4	88
Digital Signal and Image Processing	9	117
Computational Intelligence	4	35

touchscreen is mostly used for the live presentation, with a very limited amount of preprepared material.

Figure 9 presents the students' evaluation of DSP courses and a comparison of the level for standard face-to-face courses versus those presented through videoconferencing systems in 2020. The results show that the quality of the distant educational process can be at a similar level to that of traditional courses.

In particular, there is a positive experience in the students' final colloquia devoted to the discussion of selected mathematical topics or application of DSP methods. This forms another approach to share original ideas in an engaging way and to discuss the theory and implementation of modern signal processing tools, machine learning, and adaptive methods. DSP techniques are applied to data recorded by the students or obtained from specialized institutes. These discussions were initially organized in the classical face-to-face form for groups of fewer than 20 students. This model was changed recently to a virtual form, and five such videoconferencing meetings [11] were organized during 2020 with a similar number of participants during each colloquium. Students appreciated the possibility of the mutual discussion on the same level as during personal meetings. Moreover, no traveling was necessary for students from foreign countries.

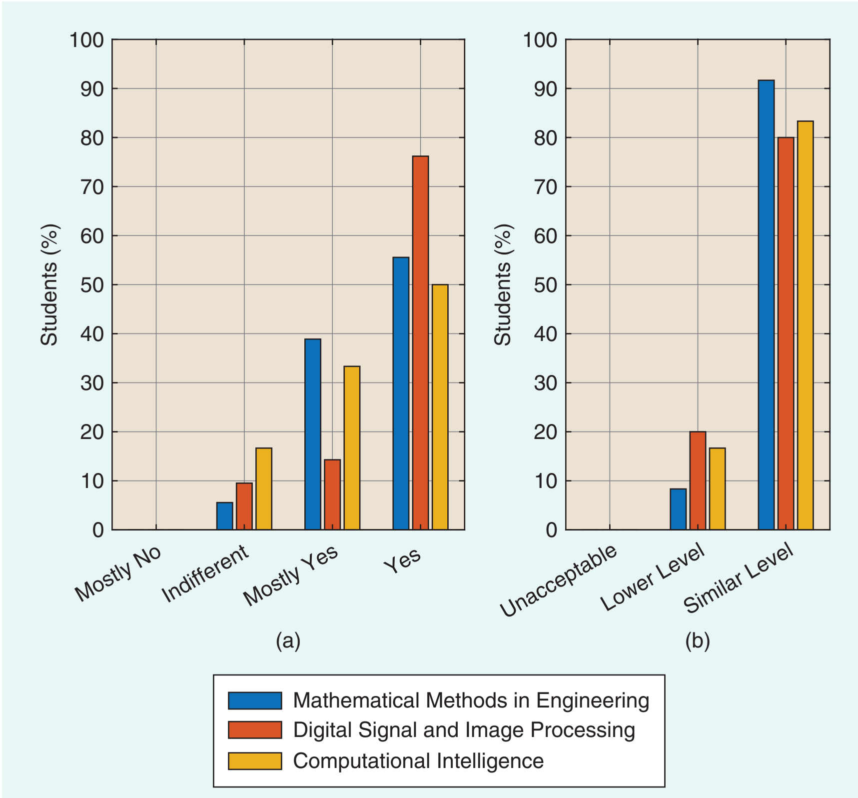

FIGURE 9. Evaluation of DSP courses by their participants, presenting (a) the level of satisfaction with different courses and (b) the comparison of distant and standard courses presented by videoconferencing systems during the year 2020.

The experience with educational courses points to the need for the connection of specialists in different areas through a suite of coherent, real-world examples. DSP can follow this direction, and it can form a very efficient platform with newly emerging educational technologies.

Conclusions

The discipline of DSP has become increasingly important in recent years across multiple fields of research, particularly for its unifying role in providing a common set of methodologies and tools with the same underlying mathematics. By applying this common approach to multiple branches of science, issues associated with overspecialization within and differentiation among scientific disciplines can be more easily resolved. It also contributes to a closer integration among sciences, as the same specialist skills are deployed across the separate disciplines. Consequently, it can encourage a closer collaboration among specialists in modern data science, academics, and students.

Education in the discipline of DSP offers the opportunity to mix a knowledge of physics, mathematics, and other areas with practical implementations in real-life applications. Modern educational courses are based on the diversity of information and the rapidly growing developments of new technologies and computing power. The interconnection of general signal processing methods and specific case studies motivates new educational technologies in studies of DSP methods. We believe in the expanding role of information engineering and signal processing as an integrating platform for different multidisciplinary topics that can reduce the problem of the "too-narrow specialization" of some scientific areas. These ideas of encouraging closer interconnection among different branches of science are similar to those of "substantial unity" suggested by Leibniz [3] and other scientists throughout history.

The current technological progress allows the introduction of new educational methods that are attractive to both students and academics. It seems that the vision of IEEE to organize future scientific conferences in both traditional and virtual ways will be very appealing to students. Moreover, virtual discussions and examinations can contribute to closer collaboration among different research groups and support the enhancement of students' creativity as well.

Acknowledgments

Thanks are due to all of the undergraduate and research students who contributed to the development and improvements of the subjects of the first author during his 50 years of teaching experience. These activities were also supported by the IEEE Society and grant agencies that enabled the international collaboration and study stays based on Erasmus, Athens, and other programs.

Authors

Aleš Procházka (a.prochazka@ieee.org) received his Ph.D. degree in 1983 and was appointed as a professor in Technical Cybernetics by the Czech Technical University in 2000. Since 1980, he has been the head of the Digital Signal and Image Processing Research Group with a focus on computational intelligence, mathematical methods of multidimensional and multichannel the nel data analysis, image processing, spatial modelling, and machine learning. He is with the Department of Computing and Control Engineering, University of Chemistry and Technology, Prague, 16628, Czech Republic and the Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague, 16628, Czech Republic. His further research interests include segmentation, feature extraction, and classification in biomedical and engineering applications. He is a member of the Institution of Engineering and Technology and a Life Member of IEEE.

Oldřich Vyšata (oldrich.vysata@fnhk.cz) received his M.D. degree in 1985 from the Medical Faculty of the University of J.E. Purkyně and his Ph.D. degree in technical cybernetics in 2011. He is with the Department of Neurology, Faculty of Medicine in Hradec Králové, Charles University, Prague, 50005, Czech Republic. His research interests focus on computational medicine, diagnostics of neurological disorders, digital signal and image processing, artificial intelligence, brain activity modelling, and machine learning. He is a member of the European Neurological Society, the Czech Society of Clinical Neurophysiology, the Czech Medical Association of J.E. Purkyně, and a Member of IEEE.

Vladimír Mařík (vladimir.marik@cvut.cz) received his M. Sc. and Ph.D. degrees in control engineering from the Czech Technical University in Prague. He is a scientific director of the Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University, Prague, 16000, Czech Republic, and the head of the Intelligent Systems Research Group with a focus on artificial intelligence, machine learning, and distributed multiagent systems to solve complex problems. His research interests include soft-computing, dynamic optimizations, environmental informatics, and computer-integrated manufacturing applications. He is a Fellow of IEEE.

References

[1] F. Erden, S. Velipasalar, A. Alkar, and A. Cetin, “Sensors in assisted living: A survey of signal and image processing methods,” *IEEE Signal Process. Mag.*, vol. 33, no. 2, pp. 36–44, 2016. doi: 10.1109/MSP.2015.2489978.

[2] S. Sanei, D. Jarchi, and A. Constantinides, *Body Sensor Networking, Design and Algorithms*. Hoboken, NJ: Wiley, 2020.

[3] G. Leibniz, A. Arnauld, and S. Voss, *The Leibniz–Arnauld Correspondence: With Selections from the Correspondence with Ernst, Landgrave of Hessen-Rheinfels (The Yale Leibniz Series)*. New Haven, CT: Yale Univ. Press, 2016.

[4] P. Prandoni and M. Vetterli, *Signal Processing for Communications*. Lausanne, Switzerland: EPFL Press, 2008.

[5] A. Oppenheim, “Back to the future [point of view],” *Proc. IEEE*, vol. 100, no. 9, pp. 2575–2579, 2012. doi: 10.1109/JPROC.2012.2204109.

[6] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury, “The dual-tree complex wavelet transform,” *IEEE Signal Process. Mag.*, vol. 22, no. 6, pp. 123–151, 2005. doi: 10.1109/MSP.2005.1550194.

[7] R. Bogner and A. Constantinides, *Introduction to Digital Filtering*. London: Wiley-Interscience, 1975.

[8] S. Kanna, W. Rosenberg, V. Goverdovsky, A. Constantinides, and D. Mandic, “Bringing wearable sensors into the classroom: A participatory approach,” *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 110–130, 2018. doi: 10.1109/MSP.2018.2806418.

[9] S. Godsill, P. Rayner, and O. Cappe, “Digital audio restoration,” *Applications of Digital Signal Processing to Audio and Acoustics (Engineering and Computer Science)*. Boston, MA: Springer, 2002.

[10] P. Rayner and W. Fizgerald, “The Bayesian approach to signal modelling,” in *Signal Analysis and Prediction (Applied and Numerical Harmonic Analysis)*. Boston, MA: Birkhauser, 1998.

[11] A. Prochazka, “DSP WWW page,” 2020, Accessed: Jan. 10, 2021. [Online]. http://uprt.vscht.cz/prochazka/prochaz.e.htm

[12] G. Strang, *Linear Algebra and Learning from Data*. Cambridge, MA: Wellesley-Cambridge, 2019.

[13] C. Jansson, R. Thottappillil, S. Hillman, S. Moller, K. Hari, and R. Sundaresan, “Experiments in creating online course content for signal processing education,” in *Proc. 45th Int. Conf. Acoustics, Speech, Signal Process.*, 2020, pp. 9215–9218.

[14] H. Li, Y. Kang, W. Ding, S. Yang, S. Yang, Y. Huang, and, and Z. Liu, “Multimodal learning for classroom activity detection,” in *Proc. 45th Int. Conf. Acoustics, Speech, Signal Process.*, 2020, pp. 9229–9233.

[15] J. Guerra, S. Dominguez-Cid, J. Guerrero, A. Garcia, D. Larios, and C. Leon, “Virtual laboratory for digital signal processing,” in *Proc. XIVth Technol. Appl. Electron. Teaching Conf.*, 2020, pp. 1–6.

[16] C. Moler, “Technical perspective: Not just a matrix laboratory anymore,” *Commun. ACM*, vol. 58, no. 10, p. 90, 2015. doi: 10.1145/2814849.

[17] A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa, and H. Phan, “Tensor decompositions for signal processing applications: From two-way to multiway component analysis,” *IEEE Signal Process. Mag.*, vol. 32, no. 2, pp. 145–163, 2015. doi: 10.1109/MSP.2013.2297439.

[18] D. Mandic, N. Rehman, Z. Wu, and N. Huang, “Empirical mode decomposition-based time-frequency analysis of multivariate signals,” *IEEE Signal Process. Mag.*, vol. 30, no. 6, pp. 74–86, 2013. doi: 10.1109/MSP.2013.2267931.

[19] J. Mareš, O. Vyšata, A. Procházka, and M. Vališ, “Age-dependent complex noise fluctuation in the brain,” *IOP Sci.: Physiol. Measure.*, vol. 34, no. 10, pp. 1269–1279, 2013. doi: 10.1088/0967-3334/34/10/1269.

[20] P. Besedová, O. Vyšata, R. Mazurová, J. Kopal, J. Ondráková, M. Vališ, and A. Procházka, “Classification of brain activities during language and music perception,” *Signal Image Video Process.*, vol. 13, no. 8, pp. 1559–1567, 2019. doi: 10.1007/s11760-019-01505-5.

[21] H. Charvátová, A. Procházka, S. Vaseghi, O. Vyšata, and M. Vališ, “GPS-based analysis of physical activities using positioning and heart rate cycling data,” *Signal Image Video Process.*, vol. 11, no. 2, pp. 251–258, 2017. doi: 10.1007/s11760-016-0928-z.

[22] D. Jarchi, J. Pope, T. Lee, L. Tamjidi, A. Mirzaei, and S. Sanei, “A review on accelerometry-based gait analysis and emerging clinical applications,” *IEEE Rev. Biomed. Eng.*, vol. 11, pp. 177–194, Feb. 2018. doi: 10.1109/RBME.2018.2807182.

[23] A. Procházka, O. Vyšata, M. Vališ, O. Tˇupa, M. Schatz, and V. Marˇík, “Use of image and depth sensors of the microsoft kinect for the detection of gait disorders,” *Neural Comput. Appl.*, vol. 26, no. 7, pp. 1621–1629, 2015. doi: 10.1007/s00521-015-1827-x.

[24] E. Buckley, C. Mazza, and A. McNeill, “A systematic review of the gait characteristics associated with Cerebellar Ataxia,” *Gait Posture*, vol. 60, pp. 154–163, Feb. 2018. doi: 10.1016/j.gaitpost.2017.11.024.

[25] A. Aristidou, J. Lasenby, Y. Chrysanthou, and A. Shamir, “Inverse kinematics techniques in computer graphics: A survey,” *Comput. Graphics Forum*, vol. 37, no. 6, pp. 35–58, 2018. doi: 10.1111/cgf.13310.

[26] A. Procházka, H. Charvátová, S. Vaseghi, and O. Vyšata, “Machine learning in rehabilitation assessment for thermal and heart rate data processing,” *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 26, no. 6, pp. 1209–12141, 2018. doi: 10.1109/TNSRE.2018.2831444.

[27] D. Cardone, P. Pinti, and A. Merla, “Thermal infrared imaging-based computational psychophysiology for psychometrics,” *Comput. Math. Methods Med.*, vol. Aug. 2015, pp. 984353:1–984353:8, Aug. 2015. doi: 10.1155/2015/984353.

[28] M. Kašparová, S. Halamová, T. Dostálová, and A. Procházka, “Intra-Oral 3D scanning for the digital evaluation of dental arch parameters,” *Appl. Sci.*, vol. 8, no. 10, pp. 1838:1–1838:9, 2018. doi: 10.3390/app8101838.

[29] A. Procházka, T. Dostálová, M. Kašparová, O. Vyšata, H. Charvátová, S. Sanei, and V. Marík, “Augmented reality implementations in stomatology,” *Appl. Sci.*, vol. 9, no. 14, pp. 2929:1–2929:13, 2019. doi: 10.3390/app9142929.

[30] T. Nakamura, Y. Alqurashi, M. Morrell, and D. Mandic, “Hearables: Automatic overnight sleep monitoring with standardized in-ear EEG sensor,” *IEEE Trans. Biomed. Eng.*, vol. 67, no. 1, pp. 203–212, 2020. doi: 10.1109/TBME.2019.2911423.

[31] A. Procházka, J. Kuchynˇka, O. Vyšata, P. Cejnar, M. Vališ, and V. Marˇík, “Multiclass sleep stage analysis and adaptive pattern recognition,” *Appl. Sci.*, vol. 8, no. 5, pp. 697:1–697:14, 2018. doi: 10.3390/app8050697.

[32] J. Wang, Y. Chen, S. Hao, X. Peng, and L. Hu, “Deep learning for sensor-based activity recognition: A survey,” *Pattern Recogn. Lett.*, vol. 119, pp. 3–11, Mar. 2019. doi: 10.1016/j.patrec.2018.02.010.

[33] S. Vaseghi, *Multimedia Signal Processing: Theory and Applications in Speech, Music and Communications*. Hoboken, NJ: Wiley, 2007.

SP

Manish Narwaria

The Transition From White Box to Black Box

Challenges and opportunities in signal processing education

©SHUTTERSTOCK.COM/LUCKYSTEP

Modern engineering education is increasingly assuming an interdisciplinary character, where developments in one area almost invariably affect other areas. A prominent example is that of signal processing, which has undergone significant changes with the emergence of machine learning (ML) and deep learning (DL) in recent years. While the impact of ML/DL is clearly visible from the viewpoint of research and development as well as industrial applications, it is not immediately clear how signal processing education should evolve in terms of pedagogy and content. Hence, the main purpose of this article is to provide some insight into this aspect. In particular, we emphasize that the introduction and popularity of ML/DL, especially at the level of teaching, has provided an opportunity to bring the focus back to some of the fundamental ideas rooted in signal processing and other related fields of study.

Background

Signal processing courses taught at both the undergraduate and postgraduate levels aim at developing an understanding of the physical world via explicit mathematical analysis. This includes both basic and advanced courses. All such courses share a common thread in that each deals with some aspect of physical signals and the underlying mathematical theory of developing analytical systems/models to process such signals. However, with the advent of ML as a major thrust area both in terms of teaching and research, certain challenges and opportunities have emerged. These range from the more simplistic application- and design-oriented challenges to the more fundamental questions related to the need for signal processing as an area of study.

As a discipline, ML has a rich history, as evidenced by many seminal theoretical ideas over the past decades, e.g., refer to the formative article by Turing [1], and several practical applications. Particularly from the viewpoint of academics, ML has always been one of the fundamental topics in mathematics, computer science, and electrical engineering, typically as an advanced level topic/area. However, the introduction of powerful computing resources, availability of data, and emergence of

Digital Object Identifier 10.1109/MSP.2021.3050996
Date of current version: 28 April 2021

DL architectures have, in some ways, completely changed the landscape in terms of how ML is viewed both by academia and industry [2]. DL, in particular, has gained popularity and is widely employed across several engineering and science disciplines.

Not surprisingly, several universities (and research centers) around the world have started emphasizing the need for ML/DL by modifying the curriculum (and overall research focus), explicitly or implicitly. This includes offering exclusive undergraduate and postgraduate degree programs in this area. For instance, the Indian Institutes of Technology at Jodhpur (i.e., the author's institute) and Hyderabad have, at the time of writing, already started dedicated undergraduate programs and specializations in data science and artificial intelligence (AI), where ML/DL will be a key component. In addition, many universities globally have started to emphasize courses on ML as early as the third semester of the undergraduate level, cutting across departments. To facilitate this, the prerequisites include basic programming skills and familiarity with calculus and probability theory, which are satisfied by most undergraduate students.

> **While the impact of ML/DL is clearly visible from the viewpoint of research and development as well as industrial applications, it is not immediately clear how signal processing education should evolve in terms of pedagogy and content.**

Many students have also quickly gained significant interest in DL as they view it as a useful tool to solve industry-oriented problems and for professional growth. This is true not just for electrical and computer engineering but increasingly so in varied domains, including chemical, infrastructure, bioengineering, and so on. All of this is, of course, desirable in a rapidly changing world. At a philosophical level, we view this as a transition from a white box to a black box to imply an emphasis on implicit modeling based on data. While, in itself, we do not view the said transition as being necessarily problematic, it has had a significant impact in terms of teaching and learning at some levels. Specifically, it does appear that DL, with its powerful modeling capabilities, has reduced the focus on certain fundamental and traditional areas of study. Similar views may also be found in the context of communications [3], multimedia processing [4], computer vision [5], and more.

This naturally leads us to a few fundamental questions on how the teaching and learning of traditional signal processing, ML, and related areas are likely to be impacted. We note that there has been a lack of systematic and explicit discussion on this. Therefore, the main purpose of this article is to put forward a few hypotheses which, in our opinion, are relevant to today's educators and research practitioners in said domains. We also attempt to provide some evidence against the said hypotheses by drawing from research and teaching experience over the past few years in the broad areas of signal processing and ML.

The emergence of the black-box approach

The need for developing interpretable/explainable ML models is probably as old as the fields of AI and ML/DL [6]. However, many ML methods, including DL, can be treated as black boxes. Also note that strictly speaking, the terms explainability and interpretability are distinct. The former refers to the post hoc analysis of a black model. The latter aims to develop white-box-type models from the beginning itself. We, however, do away with such distinction and use both terms interchangeably. However, many ML methods, including DL, can be treated as black boxes.

In simple terms, the black-box model is one where the trained weights may not have an explicit one-to-one mapping to the actual input. However, in the case of most ML methods other than DL, the said input is handcrafted (e.g., signal features are derived based on the domain and explicit techniques like transformations). Hence, in this case, at least the input to ML is not part of the said black box (i.e., there is explicit feature extraction). Consequently, input features are more amenable to scrutiny. By contrast, DL uses the data/signal as input. This means that the task of feature extraction is delegated to DL (i.e., there is implicit feature extraction). As a result, analytic or empirical analysis of the resultant DL model tends to be more difficult in practice.

System-centric or DL approach: An instructor's dilemma

Most courses in signal processing and other related areas follow what we refer to as the *system-centric approach*. This involves, in most cases, the analysis and development of mathematically tractable and transparent systems to process signals/data and examine the effects of external phenomena (e.g., the analysis of a noisy communication channel). In contrast, the DL approach emphasizes implicit modeling. For a convenient visual comparison, we show in Figure 1 the essential blocks of the two mentioned approaches: the traditional system-centric method, depicted in the flow diagram in Figure 1(a), and the DL-based method, shown in the flow diagram in Figure 1(b). We make a few important observations from this.

First, the traditional system-centric approach follows a bottom-up strategy. It seeks to build upon the knowledge that students gain from foundational and other program-specific core courses. These courses also typically emphasize a system-based approach. It may also include case studies and applications where the students may learn how to apply the learned concepts to model and/or express domain (application-specific) knowledge. All of these are expected to make the students capable of dealing with real-world applications/problems. The use of double-sided arrows in the flow diagram is intentional and emphasizes the interdependence between applications and domain knowledge expressed via basic systems concepts (e.g., distributional assumptions or the use of linear models to approximate complex systems).

In contrast, the DL-based approach depicted in the flow diagram in Figure 1(b) is relatively more top down because it is application driven. Hence, the applications primarily drive what we refer to as implicit modeling and data structure discovery. It will obviously be defined by the target of interest (e.g., supervised or unsupervised classification). From the viewpoint of implementation, a set of core activities (primarily consisting of data set creation/labeling, implementation of DL architectures via programming, and subsequent model training) then need to be carried out. As highlighted, this will require the knowledge and application of concepts mainly in calculus, linear algebra, probability theory, and optimization. However, unlike the system-centric paradigm, there is more emphasis on data due to the implicit nature of modeling that is involved. In the process, system-based prior knowledge (which, in some sense, constitutes the core of the traditional system-centric paradigm) is exploited only in a limited fashion.

Thus, the rapid emergence of ML/DL has given rise to a more hands-on and data-centric approach to problem solving [as depicted in the flow diagram in Figure 1(b)]. In light of this, it is natural for the students to explicitly (or, in many cases, implicitly) question the utility of system-centric knowledge when the relatively easier option of DL is available. This dilemma of possibly selecting one approach at the cost of the other in many ways defines a major challenge in modern signal processing education. It becomes even more prominent in cases where a real-world problem can possibly be solved using both of the approaches. Thus, it is reasonable to conclude that explicitly addressing the said challenge will help in providing much-needed perspective and context to the students.

Why the DL approach is becoming popular

The DL approach has achieved remarkable performance improvements over state-of-the-art methods in certain applications. While the ambit is increasing, the said applications still constitute a relatively small subset of possible application scenarios. Moreover, one needs to be careful in extrapolating and generalizing results from test data sets alone [7]. Still, the significant research and industry interests in DL are understandable. However, their popularity among a large section of students cutting across departments/specializations is an interesting phenomenon. We list a few plausible reasons for this.

1) As the DL framework is general, it can, in principle, be applied to several application areas for building predictive models. This requires the appropriate data (this is greatly eased due to the free availability of data sets from several application domains) and fine tuning of the DL architectures. In some ways, this allows students to at least attempt problems from domains for which they may not have any domain knowledge. In contrast, the traditional system-centric approach, which generally builds upon analytical tools, might appear to be more rigorous and limited in some sense.
2) DL has shown promise in many application areas—especially in visual signal analysis [8]. Moreover, the ML/DL community has been very active in engaging both academia and industry. (For instance, several competitions and challenges are held each year that witness high levels of participation.)
3) The extensive availability of software programs greatly facilitates students (from the beginner to advanced level) in applying DL. In fact, the convenience levels have almost reached levels that require as much as a "drag and drop" operation to build customized DL models.
4) As highlighted in Figure 1(b), the DL approach focuses more on data, implementation, and model fine-tuning. As a result, students who might not have sufficient knowledge or may

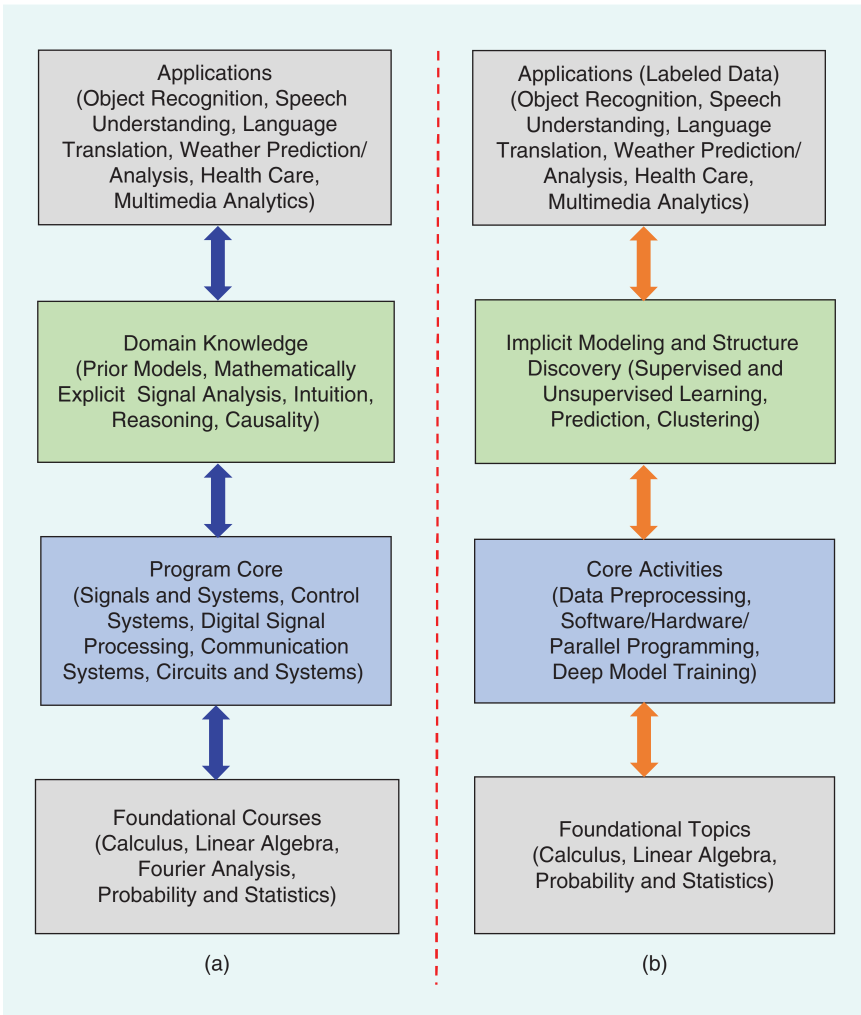

FIGURE 1. A comparison of (a) the traditional bottom-up system-centric approach and (b) the more recent top-down DL paradigm.

not have even formally studied some of the required theoretical concepts (such as convolution, optimization, and so on) find it easy to build DL models. In fact, it might not be completely wrong that say that programming skills tend to be taken as a substitute for basic concepts.
5) With the DL approach, one can visualize the data [9], preprocessing, and training process through user-friendly and convenient graphical interfaces.
6) With its ever-increasing popularity and wide acceptance across the industry [2], ML/DL is perceived by many students as a gateway to exciting career opportunities and professional growth.

The hypotheses

While one may agree or disagree with some (or all) of the reasons listed, there appears to be a consensus that the DL approach has indeed attracted tremendous interest not just from students but also from educators, researchers, and industry. It is, therefore, natural to wonder how signal processing education should evolve. To that end, we suggest the following hypotheses:

- *H1*. Most signal processing courses tend to focus on explicit signals and systems analysis. But when sufficient training data is available to enable DL-based modeling, the traditional system-centric approach does not add much value.
- *H2*. The traditional and DL-based paradigms are completely independent of each other, and selecting one over the other is essentially a zero-sum game.
- *H3*. With DL, one simply needs an application-specific labeled data set to develop an appropriate DL architecture. In that case, domain knowledge/expertise expressed via explicit signals and systems is not needed. Instead, one can simply focus on more data, faster optimization, and more efficient implementations.

In the following sections, we provide a discussion in the context of the afore mentioned hypotheses.

Why signal processing concepts are necessary

Data acquisition (including labeling in the case of supervised learning) is an integral part of an ML/DL pipeline toward model development and validation. As it relies heavily on concepts such as sampling, aliasing, quantization, clipping, device characterization, and related theoretical analysis, the importance of fundamental signal processing concepts is obvious. However, the system-centric approach (with a focus on signal processing) can also be of immense significance in providing useful technical insights into several real-world problems. These, in turn, can be exploited for developing a more explainable/interpretable ML/DL model.

Inference and prediction: An example

Most ML/DL models are developed for the task of prediction, i.e., providing a label for an unseen (or future) data point. The aspect of inference (i.e., understanding the factors/causes and their contributions) is considered, albeit implicitly. This is where the system-centric approach is useful as it allows one to make useful inferences and, in turn, predictions. We note that the distinction between inference and prediction is subtle but extremely important to understand failure cases and their severity. In this context, we note that since DL emphasizes implicit modeling, it becomes challenging to understand why it failed to provide a correct prediction in a given test condition. On the other hand, a more transparent model based on a system-centric approach will be more amenable to the scrutiny of its weakness. We emphasize here that it is not just about an accurate (or inaccurate) prediction. Rather, it is also crucial to get some solid technical insights into the strengths and weaknesses of a prediction model. This will eventually help in understanding and possibly avoiding an undesirable scenario where the trained model might learn/use counterintuitive or even unnecessary factors.

As an illustration, we consider the example in Figure 2(a) and (b), where two images are shown. These images were compressed via the JPEG 2000 algorithm at the same compression factor. They were subjectively rated on a scale from 0 (worst quality) to 5 (best quality) for their visual quality by about 30 human subjects, and the average score for each image is given below it. These images are taken from the publicly available Categorical Subjective Image Quality data set [10].

Suppose that our goal is to develop a computational model that can predict the quality of a given image. To that end, the DL approach will involve training a deep neural network architecture with the given data set, i.e., the input will be the color image, and the training target will be the average quality score. The trained model is then expected to predict the quality accurately for a given compressed test image. While this approach seems straightforward, there are two main issues. First, the trained DL model cannot explain why the image in Figure 2(a) was rated poor in comparison to the image in Figure 2(b), despite the fact that the same compression factor was used for both. Second, because it is not possible in practice to test the DL system on all of the possible test images, the generalization of its performance is questionable (even if it were to perform very well on a limited set of test images).

Alternatively, we can apply the system-centric approach. In this particular example, this would mean analyzing possible factors that contribute to visual quality and quantifying them as accurately as possible (mainly via signal processing techniques). In turn, this requires some domain knowledge. A possible solution is to consider, say, four factors, namely contrast, color, naturalness, and sharpness, all of which are well-known visual factors. These are shown below each image in Figure 2. One can observe that, for the image in Figure 2(b), all of the factors except naturalness are higher and therefore indicate a better quality for this image (which is, of course, validated from the given subjective quality scores). This also provides some insights into why human subjects rated the image in Figure 2(a) as poorer in comparison to the one in Figure 2(b). These four factors can further be combined into a single score by either using a predefined relationship or by learning it from the given data.

Note that, similar to the DL approach, it is again difficult to quantify the generalization ability in this case. However, at least the contributing factors employed are open to scrutiny both in terms of principle (i.e., whether a factor should be excluded or another one added) and their computation (as these are based on transparent signal processing techniques and approximate models of some aspects of the human visual system). Consequently, failure cases can be better analyzed.

In essence, the system-centric approach allows a case-by-case analysis if required and does not overemphasize the overall prediction accuracy alone (as is the case with the DL approach). Further, in this example, it is also possible to analyze the compression technique (in this case, JPEG 2000) for its deficiency (or advantages) vis-à-vis the considered factors. This leads to the possibility of generating new knowledge or augmenting the existing findings in the said domain. Indeed, it is not surprising that advances in several areas, including next-generation video technologies, have been enabled due to the fundamental principles and limits of human vision, which are invariably understood in terms of and quantified via signal processing techniques (e.g., contrast sensitivity function, masking effects, bandpass filtering, receptor cell response, and so on) [11]. The DL approach, on the other hand, does not provide any such insights on its own.

Toward explainable ML/DL

DL owes its success to implicit but powerful modeling capabilities. However, the implicit nature of modeling can also become a drawback in many applications due to two reasons: 1) it is not straightforward to explain why a certain prediction (right or wrong) was made by the trained DL model, and 2) the robustness of DL to adversarial examples is questionable. As a result, the development of robust and explainable ML/DL is an active research area [12]. In this context, we note that the fundamental concepts and analytical tools rooted in signal processing and related areas can play an important role. For instance, the fundamental ideas of frequency (e.g., Fourier transform) can be used to explain the generalization ability of convolutional neural networks (CNNs) [13].

Further, the authors in [14] exploit Fourier transform-based priors to improve the interpretability of DL for genomics (this involves penalizing the high-frequency components of the Fourier spectrum of input-level attribution scores). The work in [15] utilizes the idea of Fourier feature mapping to overcome the weakness of standard multilayer perceptron in representing the high-frequency content of natural images and scenes. Likewise, wavelets were exploited to develop an interpretable and frequency-aware DL model for time series analysis [16]. In a similar vein, the authors in [17] explore a wavelet-based deconvolution layer to develop a more interpretable CNN-based model for time series classification.

In terms of more recent ideas, we note that graph signal processing has been used for developing new ML algorithms with a focus on model interpretablity/explainability [18], [19]. Other signal decomposition tools have also been employed toward explainable ML/DL. For instance, the approach described in [20] exploits the well-known principal component analysis for the analysis and visualization of the learning process in DL layers.

We also reiterate that the use of contrast, color, naturalness, and sharpness as features in the context of the problem stated in Figure 2 is also an example of developing an interpretable model. Thus, as briefly discussed, a sound knowledge of signal processing and allied areas will equip students and educators with powerful concepts that can enable them to design and implement more transparent ML/DL systems.

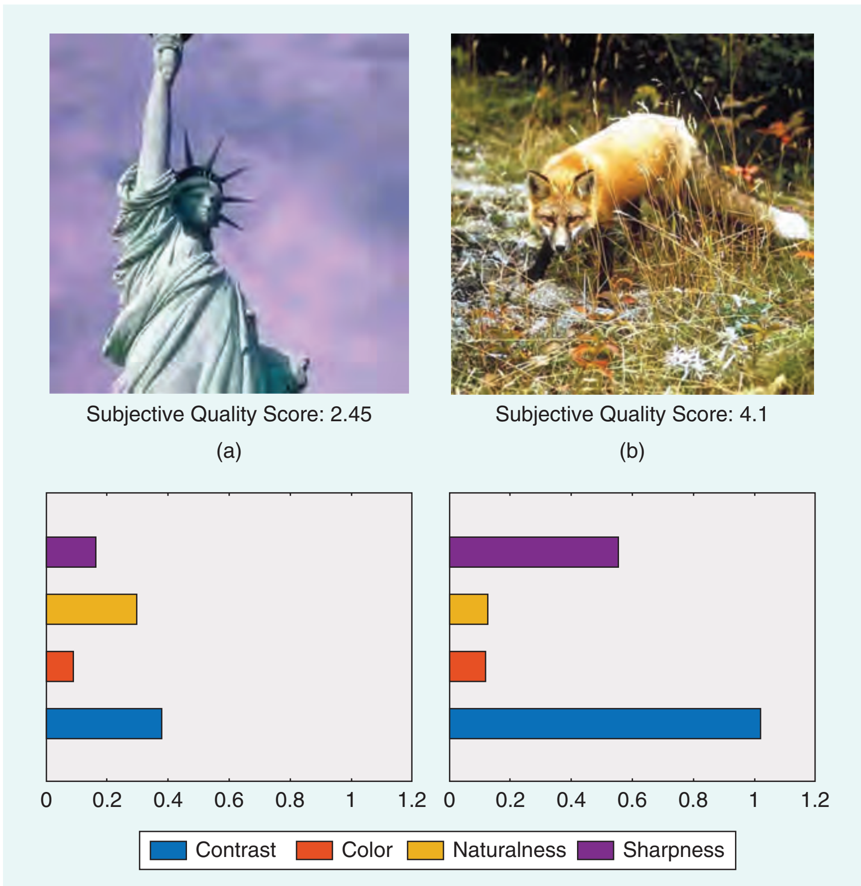

FIGURE 2. An example of image quality prediction and the importance of domain knowledge expressed via signal processing concepts. (a) A lower-quality image versus (b) a higher-quality image based on four visual factors. Several such convenient examples can be taken for the purpose, depending on the instructor interest and background of students. (Source: CSIQ dataset [10].)

Exploring commonality in signal processing and ML

Many instructors and readers will agree that teaching a signal processing course has become increasingly challenging due to increased student diversity in terms of background and the large number of courses that students are typically enrolled in. Moreover, it is likely that students find a theoretical signal processing class less attractive than, say, a course on ML/DL, AI, or computer graphics. Naturally, educators in the past have explored ways to tackle this challenge. Some examples include bringing more hands-on experience inside a signal processing classroom by using wearable sensors [21], turning the classroom into a virtual crime scene for image forensics [22], providing intuitive understanding and exploring the links between fundamental concepts [23], and so on. In this context, another complementary and reasonable strategy would be to explore common theoretical concepts and practical applications across courses. From that viewpoint, we are interested primarily in the domain of signal processing and ML/DL.

Connecting signal processing and ML through basis functions

The concept of basis functions is fundamental in signal processing. For simplicity in notation, we consider a discrete 1D signal $x[n], n \in \mathbb{Z}$. A useful and intuitive approach to analyze $x[n]$ is to decompose it in terms of basis functions. Mathematically, this is represented as

$$x[n] = \sum_k \alpha_k b_k[n], \tag{1}$$

where α_k is the coefficient of expansion and $b_k[n]$ is the basis function. Different choices for the basis functions are possible, depending on the desired properties. We have $b_k[n] = e^{j2\pi kn/N}$ [in the case of an N-point discrete Fourier transform (DFT)] or $b_k[n] = e^{j\omega n}$ (with ω being the discrete frequency of the discrete-time Fourier transform).

In the case of wavelets, $b_k[n]$ will be time limited. Thus, determining a suitable basis function is the key to analyzing $x[n]$. We note that this idea is also extended in ML for learning the decision function $f(x)$. For instance, in kernel methods [24], we have

$$f(x) = \sum_i \alpha_i \langle \phi(s_i), \phi(x) \rangle = \sum_i \langle \alpha_i \phi(s_i), \phi(x) \rangle = \sum_i \beta_i \phi(x), \tag{2}$$

where β_i is the appropriate weight/coefficient, and ϕ denotes the mapping function (possibly nonlinear). We notice that (1) and (2) are similar in that both use a linear combination of basis functions. The corresponding basis functions $b_k[n]$ and $\phi(x)$ are determined analytically (i.e., expressed mathematically). We note that $b_k[n]$ is independent of the data $x[n]$ in the case of classical Fourier or wavelet transform. That is, these functions are defined a priori, convey clear physical meaning, and are defined analytically.

In the case of graph Fourier transform, $b_k[n]$, which represents the eigenvectors of the graph Laplacian, are derived from the data and have a Fourier-like interpretation. The basis function $\phi(x)$, on the other hand, depends on the chosen kernel (e.g., Gaussian, polynomial, and so on) but is still explicitly defined. The DL approach takes this idea one step further and aims to learn the basis functions entirely from the training data without any assumptions on the functional form. As a result, unlike (1) or (2), an explicit or analytical analysis of the basis functions is not possible. Instead, only certain qualitative and data-specific arguments can be made. For example, in the case of CNNs, the learned feature maps (which may be loosely viewed as basis functions) may be visualized by approximating the reverse operations [9]. Even though the feature maps are not completely random, they are, by and large, difficult to analyze.

Thus, the definition and computation of basis functions (a concept thoroughly defined in linear algebra) provide a convenient tool to link signal processing and ML/DL and should be exploited carefully in teaching such courses. In particular, a more thorough discussion of the motivation, formulation, and benefits of different basis functions (including the ones from DL) would be helpful to provide students a context. This may help them select the more appropriate analysis tool rather than blindly apply ML/DL. Simple computer simulations that help to visualize different basis functions will also make a signal processing class more interactive and interesting for the students to grasp the central ideas.

Emphasizing signal processing concepts in ML/DL

Signal processing ideas and techniques also provide strong theoretical and practical support in terms of implementing and improving the DL framework. For example, CNNs rely on convolution, which is one of the central ideas in signal processing. Further, the pooling operation in CNNs, which aims at a reduction in the number of parameters to be learned, can also be analyzed and compared in terms of filtering operation. An instance of this is provided in Figure 3 (question 1), which the author has used in a digital signal processing class and aims to compare the average and max pooling operations. By viewing these as filters, the students are asked to explicitly provide the corresponding filter coefficients instead of using the said operations blindly in a CNN implementation. In our opinion, such a discussion can help the students to think in terms of physical interpretation and enable them to innovate from the viewpoint of real algorithmic changes instead of focusing only on improving prediction accuracy.

Moreover, recent attempts to extend DL to data defined on graphs rely on the theoretical concepts from graph signal processing, e.g., extending CNN components to graphs [25]. Another area where signal processing has historically matured and may potentially benefit ML/DL is that of online or incremental learning. This may, for instance, exploit the well-known recursive least squares algorithm [26].

ML/DL is also widely explored in many signal processing applications, e.g., a variety of visual understanding tasks [8]. However, in many cases, students tend to reduce this to a data fitting exercise, i.e., the entire focus shifts to fine tuning the DL architecture in search of better prediction accuracy. A possible solution to this can be along similar lines as those exemplified in the second question of Figure 3. In this example, the students have to make a choice between either DFT-based filtering or the use of DL to remove noise components. Surprisingly, a large percentage of students chose the DL approach despite the fact that a corresponding noise-free target signal may not even be available in this case (as a result, a supervised DL approach cannot be applied).

Another convenient and interactive example can include test data that is, for instance, under sampled. Consequently, aliasing artifacts will significantly hinder the accuracy of an ML/DL system that was trained on well-sampled data. The keen reader will agree that such examples can enable the instructor to initiate a more focused discussion about the data, its deficiencies (if any), and possible ways for mitigation, at least from the viewpoint of signal processing. As an added advantage, the analysis and performance of the trained ML/DL could possibly be improved. Thus, sound knowledge of the concepts in signal processing and related areas will enable the students to take a more holistic approach (both in terms of the data and model development) toward applying ML/DL.

Of course, more realistic and possibly customized examples from applied signal processing and other domains [like robotics, the industrial Internet of Things (IoT), and so on] can be taken, depending on the instructor and focus of the course. Finally, it is also worth pointing out that the fundamental ideas from signal processing can help in the more efficient implementation of DL. An example is that of Fourier CNN [27], which exploits the fundamental convolution theorem and fast Fourier transform to provide improvements in terms of training speed. The aspects of faster and more efficient computing are particularly important in the context of resource-constrained AI, edge computing in the industrial IoT, and embedded ML/DL.

The case in favor of domain knowledge

As alluded to in *H3*, a consequence of DL-based modeling is the potential to make domain (or expert) knowledge irrelevant to a large extent. The genesis of this lies in the fact that DL aims to learn exclusively from the data. Thus, given enough data and computational resources, one should be able to possibly solve a large variety of problems without an explicit and detailed understanding of the underlying factors and their complex interactions. Fortunately (or unfortunately), this is not quite true in several applications, ranging from health care to multimedia analytics and communication, where domain knowledge is essential.

In light of this, it is natural to wonder why domain-agnostic modeling such as that enabled via DL appears particularly attractive to many students. We have already attempted to answer this in the "Why the DL Approach Is Becoming Popular" section. As discussed, the acquisition of domain knowledge and appreciation of complex details in a sufficiently challenging application requires more effort than just a plug-and-play approach.

The emergence of domain-agnostic ML models obviously helps students simply sidestep domain knowledge yet still develop at least an initial model to solve the problem. As an undesired consequence, it is likely that students might feel that domain knowledge (expressed and analyzed in terms of signal processing) is less relevant and, hence, the overall learning process can be negatively affected. We, therefore, argue that it is crucial to anticipate such consequences and take proactive action. The simplest strategy is to discuss easy-to-grasp examples, and there can be plenty in the realm of signal processing.

1) Downsampling is a useful operation in several applications including the convolutional neural network (CNN).
We consider two ways for downsampling. First, by using average pooling and the second via max pooling.
In both cases, we use a window of length 3 i.e. downsampling is done on each non overlapping group of 3 samples. A simple illustration is given below.

$$\begin{bmatrix}1 & 2 & 3 & 4 & 5 & 6\end{bmatrix} \xrightarrow{\text{Average pooling}} \begin{bmatrix}2 & 5\end{bmatrix}$$

$$\begin{bmatrix}1 & 2 & 3 & 4 & 5 & 6\end{bmatrix} \xrightarrow{\text{Max pooling}} \begin{bmatrix}3 & 6\end{bmatrix}$$

Viewing downsampling as a filtering operation, let the filters for the case of average and max pooling be respectively given as $H_{avg} = [? \quad ? \quad ?]$ and $H_{max} = [? \quad ? \quad ?]$.Then, specify the missing values of the filter coefficients. Which filter would you recommend assuming that you need to explain the physical significance (and not merely ad hoc empirical evidence) of the choice you made?

2) Let $x(t)$ be bandlimited to 4 KHz, and it is sampled at a rate of 20 KHz to obtain $x[n]$ (with finite length N). Suppose it is known that theprobe used to capture $x(t)$ introduces two high frequency noise components (7 KHz and 8.5 KHz) in $x(t)$, and hence $x[n]$ will also be *noisy*. Then, in order to obtain a *noise-free* signal $\tilde{x}[n]$ suppose we have two options:

(a) Use DFT-based filtering and then compute the IDFT of the filtered signal.
(b) Use a machine learning method like deep learning

Which method would you recommend and why?

FIGURE 3. The sample questions used by the author to explicitly connect traditional topics in digital signal processing to ML/DL.

An example of the effect of display

We consider a single frame F from a compressed video. Assume that F is rendered (displayed) on two different displays. The first one is a standard display (with peak luminance 200 cd/m^2), while the second has a brighter display (with peak luminance 1,000 cd/m^2), as visualized in Figure 4. A question of practical interest is as follows: will the visibility of artifacts (due to compression) and hence the visual quality be different in the case of the two displays for an average human observer, and, if so, why? The question refers to the effect of the type of display used. The answer would help in several practical ways, like designing/improving video compression algorithms, display characterization, postprocessing, and so on. Thus, the goal is to predict the artifact visibility (or equivalent visual quality) for an arbitrary frame F.

To develop an ML/DL-based model for this purpose, we first need labeled data. In this example, this can be accomplished by displaying the same video frame F on the two displays one-by-one and instructing human observers to rate the quality of the rendered frame (on some scale). This needs to be repeated for a large number and variety of frames F. While it may appear that a DL model can now be developed using this labeled data set to predict the quality in a new test frame F, there is a problem due to the fact that two different types of displays were used. This implies that frame F is actually assigned two different subjective quality scores depending on the display device.

In such a situation, one can now either train two separate DL models (one for each display device) or explicitly include the display type as an input to the DL. In either case, one needs to employ explicit modeling (mainly via signal processing techniques) to analyze how the input signal (frame F) is affected by the display characteristics (e.g., luminance clipping, color gamut mapping, and so on). Subsequently, we can convert signal F into the corresponding rendered signals F_d and F_b, for darker and brighter displays, respectively. This example illustrates and reinforces the arguments made in the "Toward Explainable ML/DL" and "Emphasizing Signal Processing Concepts in ML/DL" sections, namely that solid domain knowledge based on an understanding of the data and its technical characteristics is essential for the successful deployment of ML/DL.

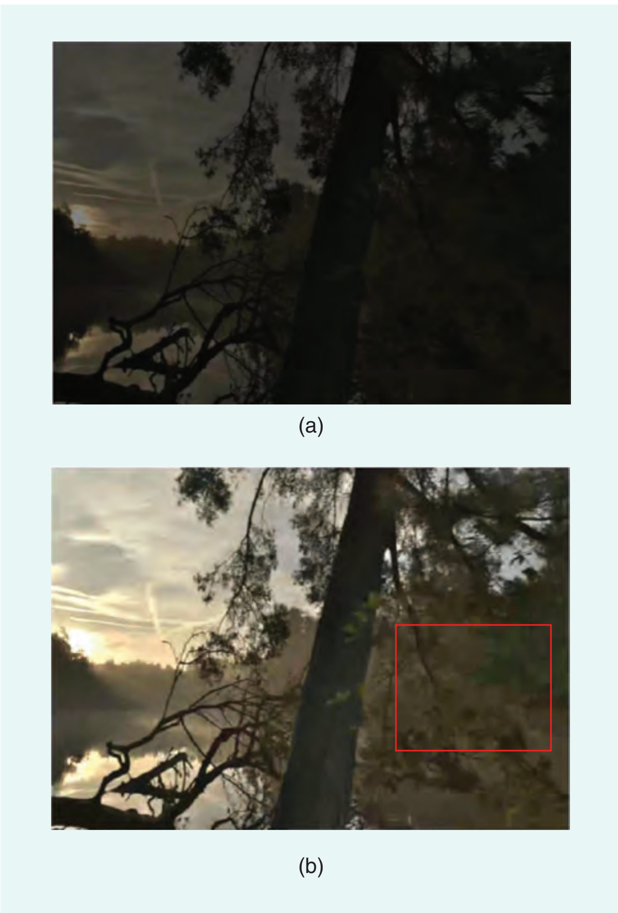

FIGURE 4. The same video frame F rendered on (a) a standard display of 200 cd/m^2 and (b) a modern bright display of 1,000 cd/m^2 will lead to different artifact visibility. Frame F is taken from a high dynamic range video data set [28]. The red square in (b) highlights an area where distortion visibility is higher.

The case of image denoising

The second example is that of image denoising. Once again, most readers will be familiar with the topic, where the goal is to reduce the noise in a given image signal. The problem has been well studied from the perspectives of signal processing and DL. Denoising via the former requires explicit modeling and certain a priori assumptions. On the other hand, a supervised DL approach trains the network with a clean signal as the target and can use the idea of residual learning [29].

To that end, one aims to minimize a loss between the estimated and target residual images. For assessing the prediction performance, measures like mean square error (MSE) or peak signal-to-noise ratio (PSNR) between the target and denoised image are widely employed. It is generally claimed that DL-based denoising outperforms traditional methods with a lower MSE (or higher PSNR). In our opinion, such conclusions tend to be simplistic in that they ignore the practical use-case scenario and instead focus only on MSE (or some similar measures).

This is where domain knowledge can play a crucial role toward more grounded comparison and validation strategies. Specifically, the purpose of denoising should first be clearly defined. If the said purpose is to generate a denoised image that is as close as possible to the clean image at the pixel level, then the use of MSE would be logical. But in the case where the goal is to produce better visually denoised images (as is the case in several image denoising applications), then an MSE-based criterion may not be suitable. In such cases, a perceptually relevant comparison will be more meaningful.

To explain this, we take the example in Figure 5(a) and (b), which illustrates an original and noisy image. In Figure 5(c) and (d), we can see that the denoised image from the CNN-based method [29] has a higher PSNR as compared to the one denoised by the well-known block-matching and 3D filtering (BM3D) algorithm [30], which employs a block-matching and collaborative filtering process and does not employ ML/DL. However, a comparison of the perceptual error maps (PEMs) [shown in Figure 5(e) and (f)] reveals that both of the denoised images appear almost the same from the viewpoint of perceptual error visibility (i.e., both PEMs indicate nearly the same probability of error visibility across the denoised images in the two cases).

Hence, one should conclude that BM3D- and CNN-based denoising methods lead to a similar performance from a perceptual viewpoint for the test image under consideration, despite the differences in PSNR values. We note that PEM, computed via the HDR-VDP-2 method [11], takes into account the viewing conditions and several aspects of the human visual system (which are expressed and analyzed in the frequency domain) and, hence, a more suitable criterion to visualize the perceptual impact of denoising.

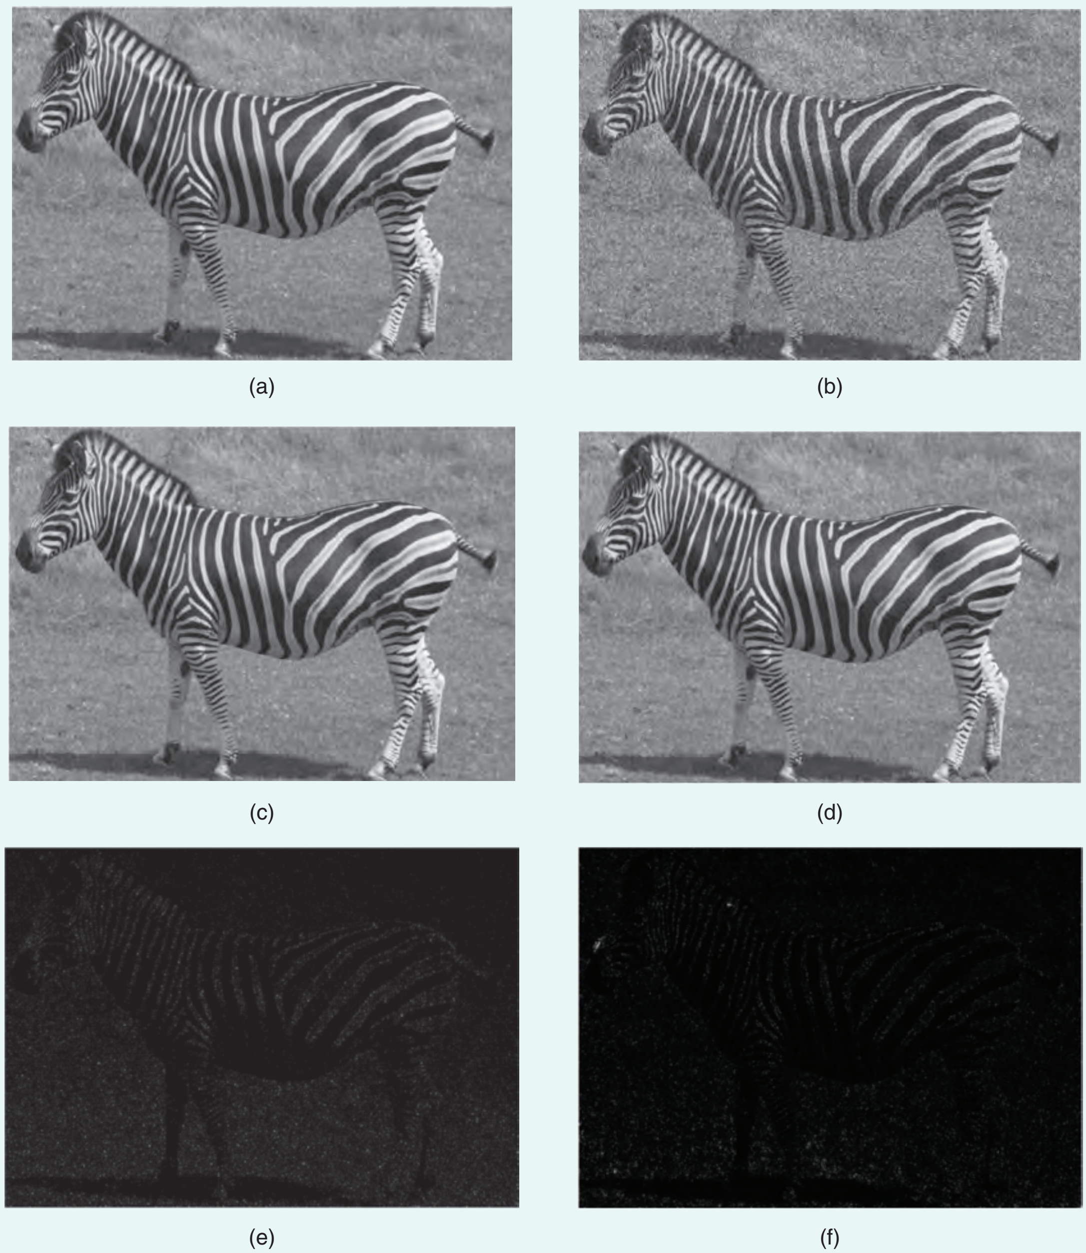

FIGURE 5. An image denoising example illustrating the importance of domain knowledge to students. (a) The original image, (b) a noisy image, (c) the denoised image using BM3D with a PSNR of 37.89 dB, (d) the denoised image using CNN with a PNSR of 38.15 dB, (e) the corresponding PEM for BM3D, and (f) the corresponding PEM for CNN. In (e) and (f), the white and black correspond to the highest and close-to-zero probability of error visibility for an average human observer.

The role of domain knowledge in model development and validation

The preceding examples were meant to illustrate the importance of domain knowledge from different dimensions. The first example (Figure 2) pertains to the extraction of relevant features (i.e., contrast, color, naturalness, and sharpness) for characterizing visual quality. It follows that domain knowledge plays a key role in developing more transparent ML systems by allowing the specification of desired features. Moreover, as already emphasized in the article, the features can be computed based on mathematical tools (related to signal processing, probability and statistics, and so on). Consequently, they are amenable to scrutiny from the viewpoint of how they are computed and if they are relevant to the problem under study. Thus, subsequent improvements from both the qualitative and quantitative perspectives are possible.

In contrast, the DL approach utilizes implicit data-dependent features. Moreover, this usually comes at the cost of fine-tuning a large number of hyper parameters. Hence, any performance improvements might be a function of the availability of higher computing power and, in turn, better tuning rather than fundamental algorithmic changes.

The second example (Figure 4) pertains to data preprocessing using knowledge of the display. The third example (Figure 5) illustrates the benefits of proper consideration of the use-case scenario (i.e., the purpose of denoising the images) and the subsequent use of perceptual measure to compare the denoising performance. As observed, this leads to different conclusions in comparison to using MSE. Other examples of exploiting domain knowledge can include the customization of the objective function (or pooling step) in ML/DL to better reflect the requirements in an application. As emphasized, all of this is enabled only through proper modeling of the domain knowledge using signal processing and related system-centric techniques.

Toward better signal processing education in the era of ML

The discussion and examples lead us to important recommendations and conclusions from the perspective of signal processing and ML education.

Recommendations

1) There is no doubt that the DL approach is generally more attractive and interesting for students. Taking inspiration from this, the pedagogical approach for signal processing and related courses can benefit immensely by following a similarly interactive and hands-on approach. Therefore, the existence of two approaches [i.e., system-centric and DL-based methods (Figure 1)] should not be viewed as a zero-sum game. Instead, explicit analysis afforded by the former should be emphasized so that students can exploit it to improve several aspects of the latter (including accuracy, explainability, robustness, incorporation of domain knowledge, architectural implementation, and so on).
2) The disadvantage of the black-box nature of the DL approach is not merely restricted to not understanding how the predictions were arrived at. Perhaps the greater cause of concern is the lack of understanding of causality of a trained DL model. As emphasized, this is not just related to overall prediction accuracy but also the severity of failure cases, where the DL model may have learned/used wrong or counterintuitive or even unnecessary factors. While more transparent modeling enabled by signal processing and related domain knowledge may not completely solve the problem of causality, it can mitigate such pitfalls. Hence, signal processing and ML courses will benefit by taking a more holistic approach and exposing the students to specific pros and cons.
3) The DL approach owes its ever-growing popularity in part to the success it has had in competitions and challenges organized at international levels (for example, the ImagNet visual recognition challenge; see http://www.image-net.org/). While there is no denying that the DL approach has achieved lower error rates than competing methods on some difficult tasks (in computer vision, speech recognition, and so on), such a winner-takes-all approach may not always be optimal, especially from the viewpoint of student learning. Thus, instead of taking a top-down and accuracy-driven approach, it is crucial that students are exposed to the nuances and first principles (and their commonalities) of signal processing and ML. We also argue that ML/DL should be treated as another tool and not the only one. Indeed, the drawbacks of both the traditional system and DL-based modeling should be emphasized, preferably through practical examples in both signal processing and ML courses. Finally, the role of domain knowledge expressed through signal processing cannot be emphasized enough. Indeed, several educators and researchers have cautioned against forgetting the basic principles of human vision, understanding of scene and image formation, causality, reasoning, and so on [4], [5].
4) Based on the discussion in this article, we argue that the importance of signal processing and related concepts is well beyond mere data acquisition. As discussed, fundamental concepts learned from the system-centric approach can help to improve ML/DL systems from the viewpoint of analysis, design, and validation. The discussion and examples thus provide evidence against hypotheses *H1*, *H2,* and *H3* formulated in the section "The Hypotheses."

Final remarks

A strong motivation for this article stems from the fact that students in today's world tend to enroll in a large number of interdisciplinary courses, possibly across departments. Accordingly, the pedagogical changes (if needed) should be bottom up, i.e., connecting theoretical concepts and ideas across courses. In that context, signal processing and ML courses lend themselves to this almost naturally, with some effort on the part of the instructor and students. To facilitate that, we suggested three hypotheses (*H1*, *H2,* and *H3*). We believe that evidence against (or in favor) these will provide a reasonable baseline to take the discussion forward on how signal processing education might evolve in an era of ML/DL.

While the arguments laid out in this article provide some evidence against the mentioned hypotheses, we do not claim to have completely rejected them. Instead, the examples and discussion were primarily geared toward identifying certain pedagogical changes that might help both students and instructors, to provide a better context for signal processing education. While the examples drew heavily from the authors' research interests and background, the focus was more on connecting the dots between signal processing and ML. Hence, as emphasized in the article, the examples and discussion could easily be tailored according to the instructor and student background. Despite the positioning of the article from a signal processing perspective, we strongly believe that the hypotheses, ideas, and discussion are equally applicable for teaching relevant courses in ML/DL.

Author

Manish Narwaria (narwaria@iitj.ac.in) received his Ph.D. degree in computer engineering from Nanyang Technological University, Singapore, in 2012. He was a researcher with the Laboratory of Digital Sciences of Nantes, Image Perception Interaction team, France, until 2016. Currently, he is an assistant professor in the Department of Electrical Engineering at the Indian Institute of Technology Jodhpur, Karwar, Rajasthan, 342037, India. His research interests include the area of multimedia signal processing with a focus on the perceptual aspects toward content capture, processing, and transmission. He is a Member of IEEE.

References

[1] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. LIX, no. 236, pp. 433–460, 1950. doi: 10.1093/mind/LIX.236.433.

[2] Q. Le, L. Miralles, S. Kulkarni, J. Su, and O. Boydell, "An overview of deep learning in industry," in *Data Analytics and AI*. Boca Raton, FL: CRC Press, July 2020, pp. 65–98.

[3] G. Frantz. "The death and possible rebirth of DSP." IEEE ComSoc Technology News (CTN). 2017. [Online]. Available: https://www.comsoc.org/publications/ctn/death-and-possible-rebirth-dsp

[4] J. Sang, J. Yu, R. Jain, R. Lienhart, P. Cui, and J. Feng, "Deep learning for multimedia: Science or technology?" in *Proc. 26th ACM Int. Conf. Multimedia*, 2018, pp. 1354–1355. doi: 10.1145/3240508.3243931.

[5] R. Chellappa, "The changing fortunes of pattern recognition and computer vision," *Image Vis. Comput.*, vol. 55, pp. 3–5, Nov. 2016. doi: 10.1016/j.imavis.2016.04.005.

[6] Y. Kodratoff, "The comprehensibility manifesto," *AI Commun.*, vol. 7, no. 2, pp. 83–85, 1994. doi: 10.3233/AIC-1994-7201.

[7] V. Vapnik, *Statistical Learning Theory*. New York: Wiley, 1998.

[8] F. Porikli, S. Shan, C. Snoek, R. Sukthankar, and X. Wang, "Deep learning for visual understanding [From the Guest Editors]," *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 24–25, 2017. doi: 10.1109/MSP.2017.2744538.

[9] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Computer Vision–ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham: Springer-Verlag, 2014, pp. 818–833.

[10] E. C. Larson and D. M. Chandler, "Most apparent distortion: Full-reference image quality assessment and the role of strategy," *J. Electron. Imag.*, vol. 19, no. 1, pp. 1–21, 2010. doi: 10.1117/1.3267105.

[11] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich, "HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions," in *Proc. ACM SIGGRAPH 2011 Papers*, 2011, pp. 1–14. doi: 10.1145/1964921.1964935.

[12] N. Xie, G. Ras, M. van Gerven, and D. Doran, "Explainable deep learning: A field guide for the uninitiated," 2020, arXiv:2004.14545.

[13] H. Wang, X. Wu, Z. Huang, and E. P. Xing, "High-frequency component helps explain the generalization of convolutional neural networks," in *Proc. 2020 IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 8681–8691. doi: 10.1109/CVPR42600.2020.00871.

[14] A. M. Tseng, A. Shrikumar, and A. Kundaje, "Fourier-transform-based attribution priors improve the interpretability and stability of deep learning models for genomics," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020.

[15] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, and R. Ng, "Fourier features let networks learn high frequency functions in low dimensional domains," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020.

[16] J. Wang, Z. Wang, J. Li, and J. Wu, "Multilevel wavelet decomposition network for interpretable time series analysis," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2018, pp. 2437–2446. doi: 10.1145/3219819.3220060.

[17] H. Khan and B. Yener, "Learning filter widths of spectral decompositions with wavelets," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, pp. 4606–4617.

[18] X. Dong, D. Thanou, L. Toni, M. Bronstein, and P. Frossard, "Graph signal processing for machine learning: A review and new perspectives," *IEEE Signal Process. Mag.*, vol. 37, no. 6, pp. 117–127, 2020. doi: 10.1109/MSP.2020.3014591.

[19] R. Anirudh, J. J. Thiagarajan, R. Sridhar, and T. Bremer, "Margin: Uncovering deep neural networks using graph signal analysis," 2018, arXiv:1711.05407.

[20] X. Li and F. Li, "Adversarial examples detection in deep networks with convolutional filter statistics," in *Proc. 2017 IEEE Int. Conf. Comput. Vis. (ICCV)*, pp. 5775–5783. doi: 10.1109/ICCV.2017.615.

[21] S. Kanna, W. von Rosenberg, V. Goverdovsky, A. G. Constantinides, and D. P. Mandic, "Bringing wearable sensors into the classroom: A participatory approach [SP Education]," *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 110–130, 2018. doi: 10.1109/MSP.2018.2806418.

[22] C. Pasquini, G. Boato, and R. Bohme, "Teaching digital signal processing with a challenge on image forensics [sp education]," *IEEE Signal Process. Mag.*, vol. 36, no. 2, pp. 101–109, 2019. doi: 10.1109/MSP.2018.2887214.

[23] D. P. Mandic, S. Kanna, and A. G. Constantinides, "On the intrinsic relationship between the least mean square and Kalman filters [lecture notes]," *IEEE Signal Process. Mag.*, vol. 32, no. 6, pp. 117–122, 2015. doi: 10.1109/MSP.2015.2461733.

[24] T. Hofmann, B. Schölkopf, and A. J. Smola, "Kernel methods in machine learning," *Ann. Statist*, vol. 36, no. 3, pp. 1171–1220, 2008. doi: 10.1214/009053607000000677.

[25] M. Cheung, J. Shi, O. Wright, L. Y. Jiang, X. Liu, and J. M. F. Moura, "Graph signal processing and deep learning: Convolution, pooling, and topology," *IEEE Signal Process. Mag.*, vol. 37, no. 6, pp. 139–149, 2020. doi: 10.1109/MSP.2020.3014594.

[26] Y. Engel, S. Mannor, and R. Meir, "The kernel recursive least-squares algorithm," *IEEE Trans. Signal Process.*, vol. 52, no. 8, pp. 2275–2285, 2004. doi: 10.1109/TSP.2004.830985.

[27] H. Pratt, B. Williams, F. Coenen, and Y. Zheng, "FCNN: Fourier convolutional neural networks," in *Machine Learning and Knowledge Discovery in Databases*, M. Ceci, J. Hollmén, L. Todorovski, C. Vens, and S. Džeroski, Eds. Cham: Springer-Verlag, 2017, pp. 786–798.

[28] J. Froehlich, S. Grandinetti, B. Eberhardt, S. Walter, A. Schilling, and H. Brendel, "Creating cinematic wide gamut HDR-video for the evaluation of tone mapping operators and HDR-displays," in *Digital Photography X*, N. Sampat, R. Tezaur, S. Battiato, and B. A. Fowler, Eds., vol. 9023, Bellingham, WA: International Society for Optics and Photonics, 2014, pp. 279–288.

[29] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, 2017. doi: 10.1109/TIP.2017.2662206.

[30] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, 2007. doi: 10.1109/TIP.2007.901238.

SP

Irwin King, Chandni Saxena, Christian Pak,
Chak-ming Lam, and Haiyan Cai

Rethinking Engineering Education

Policy, pedagogy, and assessment during crises

©SHUTTERSTOCK.COM/LUCKYSTEP

Crises, especially the recent COVID-19 pandemic, have significantly impacted traditional teaching pedagogy, which often relies on face-to-face interactions. It is crucial that various stakeholders in education, including administrators, staff members, teachers, parents, learners, government officials, and so on, adapt to abrupt changes and disruptive transformations caused by emergency situations. In this article, we map out approaches to stakeholders that underpin teaching and learning effectiveness for engineering education (EE) in terms of policy, pedagogy, and assessment. The contributions of this article are threefold. First, we revisit a framework that enables administrators to devise policies for a secure and safe learning environment. Second, we propose Crisis-Resilience Pedagogy (CRP), which highlights and integrates important attributes such as adaptability, creativity, connectivity, diversity, and endurance into pedagogical components for effective teaching and learning. Third, we outline how to leverage education technology for outcomes assessment. To illustrate the challenges, solutions, and possibilities in this "new normal," we utilize and reflect on the results of an observational study conducted during the pandemic. Our approaches can be easily extended to other academic disciplines in other institutions to strengthen the resilience of our education systems in times of crisis.

Introduction

We live in a natural world characterized by risk, catastrophe, and instability resulting from disasters and crises. These abrupt events, small or large in scale, natural or anthropogenic, have had a profound impact on individuals, organizations, communities, and states on a worldwide level. In recent years, we have observed a shift in the nature, causes, frequency, consequences, and adversities associated with crises [1], which has led to new challenges and growing research interests to understand and manage these events. In the past few decades, crises like natural disasters, armed conflicts, social movements, and pandemics have severely impacted education in different parts of the world, e.g., the 2010 floods in Pakistan, the 2013 Ebola outbreak across West Africa, the 2018 civil war in Syria, and

Digital Object Identifier 10.1109/MSP.2021.3059243
Date of current version: 28 April 2021

the Black Lives Matter movement in the United States [2], [3]. In 2020, the COVID-19 outbreak took center stage, forcing the world to face the suspension of normal academic activities, which has affected more than 1.6 billion learners in more than 190 countries around the world by August of the same year [4]. We now have entered a "new normal," where education may intertwine with sudden disruptions. It is our hope that this article provides some guidelines and reflections to assist various stakeholders in sustaining education during an unprecedented period and reaching long-term academic goals in future crises.

Without a loss of generality, we examine the key components in EE at both the micro and macro levels (see Figure 1) to ensure effective teaching and learning during emergencies, with our findings applicable to a wide range of disciplines. At the micro level, we consider these three components: knowledge (content), delivery (pedagogy), and assessment (outcome). At the macro level, we focus on policies that govern these three key factors in EE. In the "Crisis Management in EE" section, we present a crisis management framework (CMF) for education, which is useful for institutions to develop policies to provide a safe, productive, and flexible online teaching and learning environment that achieves effective educational experiences. In the "Crisis-Resilience Pedagogy" section, we share some key resilience attributes that can be incorporated into the pedagogy in the delivery of knowledge so that student learning remains uninterrupted during adversity. We describe outcomes assessment in the "Education Technology and Online Assessment" section, where the area of focus is on how to evaluate students' online learning outcomes fairly in difficult environments. The interactions among these three components under well-conceived policies facilitate effective teaching and learning during crises.

To support our findings and discussions, we use case examples and survey results of teaching and learning experiences from the Chinese University of Hong Kong (CUHK), a comprehensive and diverse education institution with more than 20,000 students and an active international and exchange student body. The university has experienced major crises as a result of social movements as well as the COVID-19 pandemic in recent years, which have led to prompt transformations of the process of teaching and learning. The outcomes of such actions were collected through course and teaching evaluation (CTE) surveys. Sharing our findings, along with the reflections of earlier-defined components with other stakeholders, provides insights that may benefit the overall planning and lead to a better education in times of crisis.

Crisis management in EE

CMF

As the year-round pandemic has created a severe disruption to global education systems, this section introduces a CMF to support teaching and learning in various academic disciplines, including EE.

A CMF is key to mitigating the traumatic effects of crises with timely action as a high percentage (65 out of 82) of crisis incidents result in positive outcomes due to crisis interventions [5]. Despite a plethora of CMFs, only a few of these are suitable for education. In [6], Smith outlines a three-stage framework, which includes a crisis of management, operation of crisis, and crisis of legitimation for industries. In the nonprofit sector, Salter presented a prevention, preparedness, response, and recovery (PPRR) framework for preserving governmental operations [7]. Although the PPRR framework was first introduced many years ago, it is still quite applicable in current-day crises. Nonetheless, we further contextualize how the use of technology and in-depth reflections can make the PPRR model more relevant, as shown in Figure 2.

Prevention

Prevention is an indication of any activity that administrators undertake to prepare for situations and considers important factors in regard to feasibility, fairness, and flexibility (3Fs). *Feasibility* is a determination of whether policies can be implemented for school members with tangible resources so as to convey teaching content and accomplish learning goals. Shifting classes to an online format, for example, may be a viable alternative to continuing education at school, given adequate Internet resources. *Fairness* pertains to the balance of policies to different stakeholders (e.g., the consistency of examination policies for local, international, and exchange students) to ensure education quality. *Flexibility* refers to providing more choices to stakeholders during crises. For instance, schools can implement more assessment and credit options, such as "Pass/Fail," "Credit/Noncredit," and so forth to assist teachers and students during strenuous situations. By devising policies with the "3Fs," administrators can prepare for unexpected predicaments.

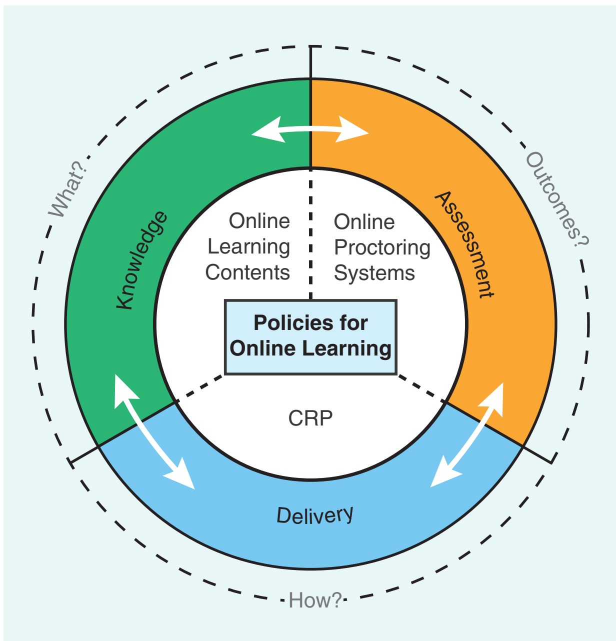

FIGURE 1. Online learning implementation and development.

Preparedness

Preparedness is a measurement of whether the stakeholders are ready to prevent additional losses, effectively through examining the scalability, rapidity, and intensity of the incident. The scalability of the event reveals the number of stakeholders that are affected by crises. Policy makers can use it to allocate and equip resources for the shifted-education mode. In addition to scalability, administrators should investigate the rapidity of crises so that timely responses can be offered to counter further deterioration. For instance, during the quickly escalating COVID-19 pandemic, the Medical School of CUHK suspended normal learning and amended curricula to address social-distancing measures by staggering lab sessions. Policy makers should also assess the intensity of crises to address the potential damage that a disruption may cause; they can use it to manage the duration of interventions. Detecting the aforementioned measurements in the preparedness stage enables policy makers to counter a crisis effectively.

Response

In this stage, immediacy, mitigation, and maintenance are the major factors that administrators should consider. *Immediacy* refers to the timeliness of deploying the prepared actions. Soon after the spread of the COVID-19 pandemic, CUHK purchased and introduced new application tools within a week and announced the arrangement of online education for all school members through the prescribed official channels. In the meantime, *mitigation* is a tactic that is used to resolve problems so that school members can adapt the transformation in an agile manner. For instance, CUHK received positive feedback from school members who used its ticketing systems for prompt responses to policy, administrative, technical, and other inquiries. *Maintenance* means that administrators should ensure that all infrastructure resources are in a functional condition to avoid additional disruptions. For instance, a timeline for regular resource checkups alongside updates, such as network resources; disk storage; remote laboratory instruments; and so on, is constructive to maintain a smooth online education. These three factors help administrators to prevent losses from predicaments.

Recovery

Although schools will resume normal operation following predicaments, it is crucial for educators to improve education through evaluation, refinement, and incorporation. *Evaluation* tabulates the effectiveness of crisis management policies through teacher/student feedback. For instance, we collected and reviewed CTE feedback regarding crisis interventions for education quality assurance. Upon assessment, administrators can refine policies based on the evaluation result to prepare for future crises. The *refinement* process improves crisis-related policies using the lessons learned. *Incorporation* involves integration of crisis interventions with regular teaching to innovate education. One suggestion is that administrators may incorporate online lessons alongside face-to-face lessons to bolster teaching and learning effectiveness under the new normal. These three actions allow administrators to improve education quality and prepare for future crises.

Reflections on crisis interventions

In this section, we offer some insights about the effectiveness of crisis intervention gained from our experiences during the pandemic.

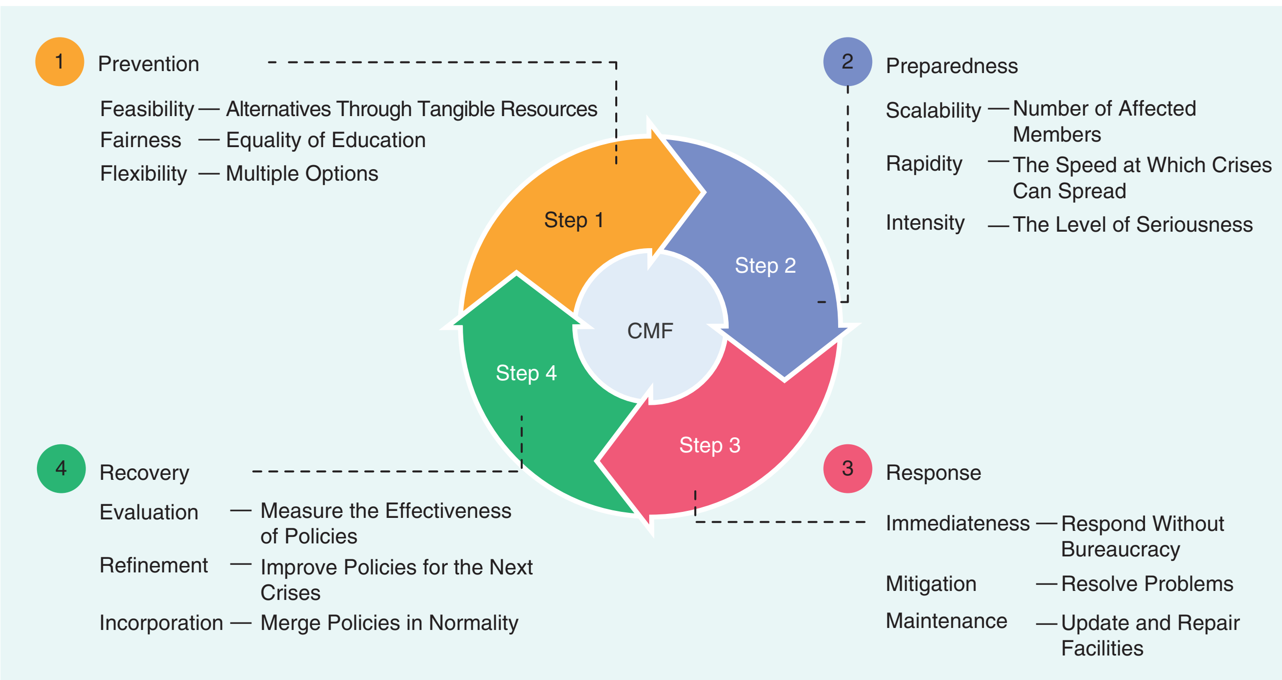

FIGURE 2. The stages of a CMF.

1) *Additional work is evident for all stakeholders*: It is evident that all stakeholders need more time for preparation, response, and reflection, but the online mode also provides some good tradeoffs for ensuring safety during the pandemic. In CTEs, students reported that they needed extra time for completing assignments to achieve remote learning outcomes. Disruptions in the original curricula have increased the workload for teachers, who must prepare additional materials for instruction. Educators at the university level agree that the assignment and preparation workload for online lessons has increased. This has also affected technical teams, which must now provide support after hours. It is essential for administrators to not underestimate these extra workloads and allocate the appropriate resources (hardware, software, manpower, and so forth) to stakeholders for addressing teaching and learning stress and acculturating to the postcrisis norm.
2) *A crisis management policy is paramount*: Having a crisis management policy is invaluable; however, swift and appropriate action is helpful to alleviate impacts. Despite difficult situations, online education is effective for students to demonstrate learning outcomes. Based on more than 900 open-ended feedback responses from engineering students, it was found that students could save time in traffic and utilize the time to better prepare for courses. This also stipulates that a shift of education mode is safe at the risk of having face-to-face teaching and learning.

To address these concerns, an effective crisis management depends on preparation, the quality of the appropriate infrastructure, and attainable manpower. A CMF that utilizes PPRR enables administrators to overcome chaotic circumstances.

CRP

In this section, we first introduce the CRP framework, which can be adopted to assist teaching and learning during crises. This pedagogy is characterized by its resilience to unprecedented changes and its applicability to different academic disciplines, including engineering. Teachers can incorporate CRP into online learning to promote the effective delivery of content to students in times of emergency. In a later part of this section, we analyze learning outcomes from the successful implementations of various courses in terms of online lessons, labs, and capstone projects. We consider the online pedagogical experience and identify some pedagogical challenges. With the persistent pandemic situation, we further suggest solutions to address these issues.

CRP model

To prepare for unexpected situations, teachers need to incorporate resilience into pedagogy and devise teaching methods that are adaptable to changes quickly. *Resilience* refers to the ability to recover rapidly from adversity; it is a continuous process of maintaining sustainability in difficult times [8]. To sustain teaching and learning during adversity, our previous work [9] proposes a CRP model that incorporates resilience into education with five attributes: 1) adaptability, 2) creativity, 3) connectivity, 4) diversity, and 5) endurance (see Figure 3). In the following sections, we describe how each attribute can be used for teaching and learning during crises.

Adaptability

Teachers need to adapt to changes in crises quickly and effectively. Because traditional ways of teaching are unavailable, they should use a flexible pedagogy to continue their teaching. Soon after the outbreak of COVID-19, teachers switched from traditional classrooms to online lessons with the use of video-conferencing tools such as Zoom and Microsoft Teams in response to campus closures. They adopted modified materials and assessment rubrics for online learning. Making adjustments and refinements to traditional pedagogy to suit the appropriateness of the situation can help students to sustain learning.

Creativity

Designing creative pedagogy using different online education platforms and tools enables students to learn effectively during adversity. Teachers are able to design interactive learning activities for students, allowing students to learn in an enjoyable way. One example of creative pedagogy is gamification, which involves the incorporation of game elements into education. Teachers can utilize online points, badges, and leaderboards to encourage competition among students, enhancing their learning motivation. Using creative elements in pedagogy can assist and promote student learning in times of emergency.

Connectivity

Teachers and students need to stay connected during crises. The connection among teachers and students should be multidirectional instead of unidirectional. Although teachers deliver teaching content to students, students should also be able

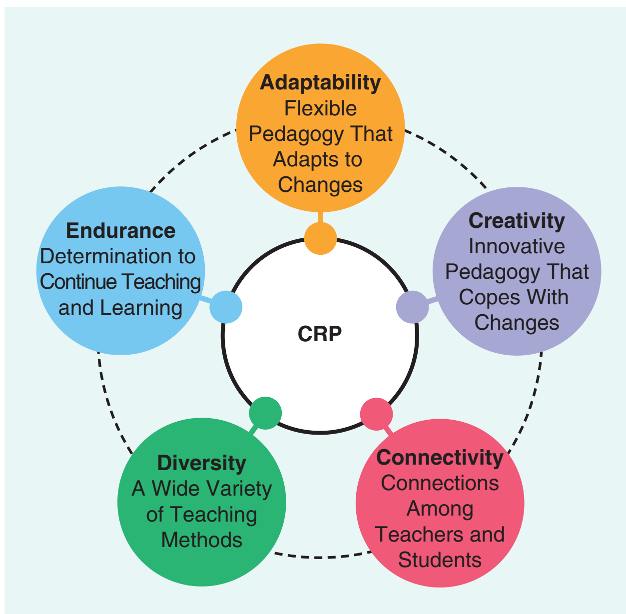

FIGURE 3. The key attributes in CRP.

to provide feedback to teachers. With technological tools like BlueJeans and Zoho Meeting, teachers can receive instant responses from students. Connections among students are also important in peer learning. Students need to connect among themselves through online collaborative platforms such as Apple Classroom, ProofHub, and MindMeister during crises so that they are able to work with each other to complete assignments and projects.

Diversity

Teachers should deploy a wide variety of teaching approaches to cater to learner diversity. Without a physical presence, students with special education needs may find it difficult to continue learning; therefore, it is crucial for teachers to devise a suitable pedagogy according to their needs. For instance, it can be difficult for students with hearing impairments to understand verbal instructions during online lessons in the absence of lipreading and assistance from classmates. In such cases, teachers may create dialogues of video lessons with automated captioning tools like Communication Access Real-Time Translation (CART) and Kapwing. In addition to online lessons, teachers can prepare other materials, such as animations, to facilitate their learning. Leveraging a set of diversified learning tools can serve the different needs of individual students during times of adversity.

Virtual laboratories, including CircuitLab, LabVIEW, and Labwork, are simulators that allow students to carry out tests such as circuit and digital signal processing experiments.

Endurance

To continue teaching and learning during adversity, teachers and students must possess *endurance*—the determination to reach goals. Students can develop this capability by setting specific and attainable learning goals, such as the number of online micromodules enrolled in per month. They can also utilize different apps to stay focused and motivated when they learn. For instance, they can use Cold Turkey to block websites that are unrelated to study. Nurturing endurance spurs students to continue to learn.

Teachers can apply these key attributes in CRP to devise suitable pedagogy to promote learning. Online learning embodies these attributes well because it is adaptable and allows students to learn anywhere and at anytime. As a wide variety of online platforms and tools are currently available, creative teaching and learning approaches can be devised more easily. Online learning can break down the barrier of time and space, connecting teachers and students. The diversity of online learning tools and methods makes learning fun and engaging, thus encouraging students to overcome difficulties and keep learning. It is essential for teachers to consider how best to utilize different online learning approaches under the framework of CRP when faced with crises.

Applications of CRP in EE

In this section, we suggest how educators can apply CRP to EE in times of emergency by referencing some successful examples from during COVID-19. A typical engineering curriculum consists of 1) lower-division courses, 2) laboratory courses, 3) project-based learning, and 4) upper-division and postgraduate-level courses. Students are introduced to rudimentary knowledge and theories in lower-division courses. In laboratory courses, they learn practical skills and acquire hands-on experience. Through carrying out projects, students learn to cooperate with others and develop problem-solving skills. Upper-division and postgraduate-level courses encourage students to acquire specialized knowledge and skills that are useful in their future careers. All of these courses are essential for nurturing students' ability to integrate knowledge and hands-on skills to be successful engineers. However, these courses are being affected differently during crises.

During emergencies, the traditional ways of conducting theoretical courses and laboratory courses become infeasible due to campus closures. Carrying out projects is also difficult because students cannot easily meet with one another. In view of these challenges, teachers can use CRP to modify teaching approaches in these courses during adversity (see Figure 4).

Lower-division courses

With the adaptability attribute, teachers can encourage students to adopt a personalized learning approach and take the initiative to learn even when face-to-face classrooms are unavailable. A personalized learning approach increases students' intrinsic motivation and autonomy in learning [10]. Students can make use of self-paced online courses such as massive open online courses (MOOCs) and small private online courses to gain relevant, basic knowledge in lower-division courses. These courses contain modularizable content small enough to be absorbed in a short period of time that can be easily configured to adapt to different situations and scenarios. Teachers can design micromodules for students so that they can take a personalized learning path by adjusting their pace of learning. For instance, our engineering faculty has produced a series of micromodules for foundation courses: "Introduction to Engineering Design" and "Engineering Physics." These micromodules have successfully encouraged personalized learning. As students spend more time at home in times of crisis, they can allocate their time to learn according to their preferences.

Teachers can even integrate self-paced online courses with the flipped-classroom approach into their teaching. This blended learning model enables students to adapt to changes and continue their learning during crises. Teachers can apply the diversity attribute when designing content for these courses by making use different multimedia to help students understand concepts and theories better. Rakhashia et al. [11] share their experience of conducting a MOOC, which involves fundamental mathematical knowledge and

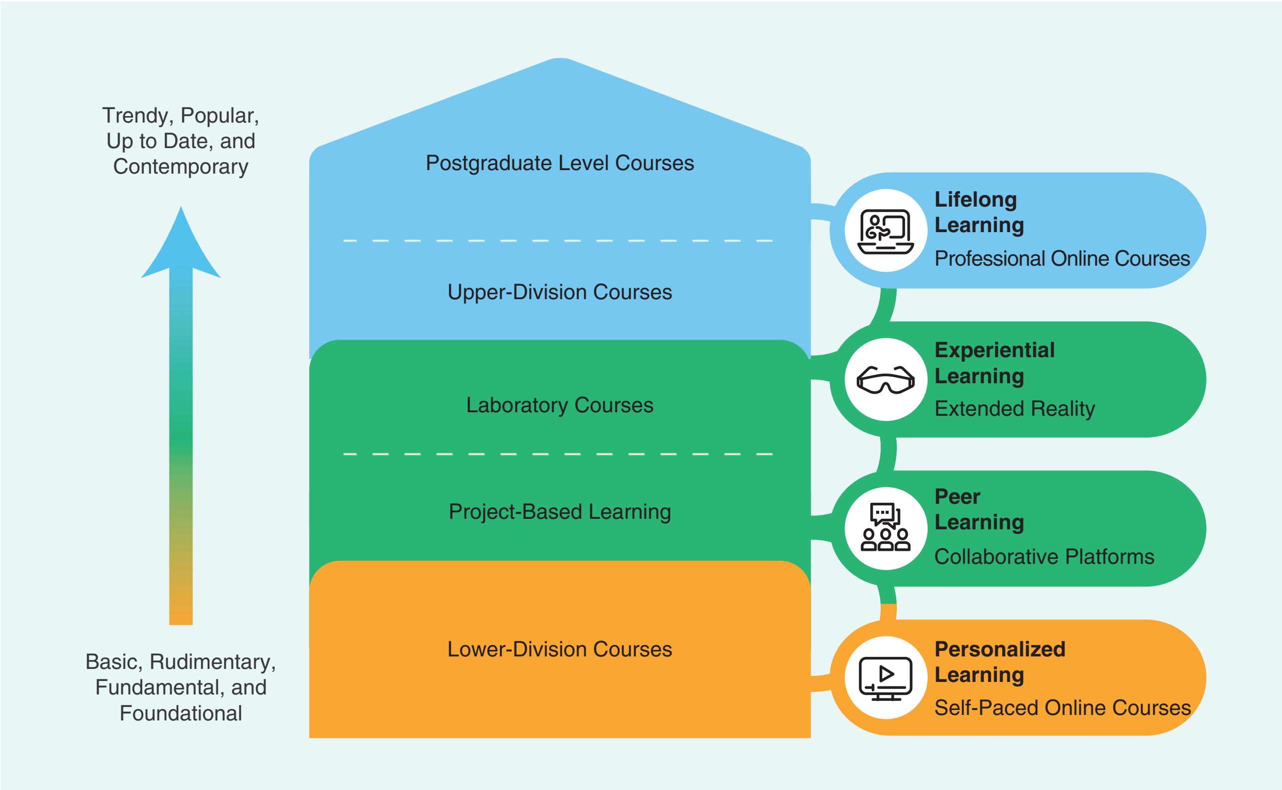

FIGURE 4. A CRP for EE.

the engineering concepts used in signal processing. To help students understand how signals are reconstructed using Fourier series components, they used programming software such as MATLAB and Python to create visuals and animations. Using multimedia in course design not only enables students to gain a better understanding of the course material but also makes learning more entertaining and engaging.

Laboratory courses

Using the creativity attribute in CRP, teachers and students can employ innovative approaches to continue laboratory courses during crises. Teachers can consider the use of extended reality (augmented/virtual/mixed/extended reality) in constructing virtual and remote laboratories to teach practical skills. Virtual laboratories, including CircuitLab, LabVIEW, and Labwork, are simulators that allow students to carry out tests such as circuit and digital signal processing experiments. Remote laboratories are tools that connect digital devices to real laboratories. Teachers and students can work together in remote laboratories even when they are physically apart. During the pandemic, instructors from CUHK creatively utilized an online robotic laboratory to teach the robotic course "Robots in Action," as shown in Figure 5. Students used their own computers at home to control the robots in the laboratory remotely, successfully gaining hands-on skills from the robotic remote laboratory. Moreover, the rate of positive feedback on the laboratory learning experience during the course was 87.1% [12]. This example shows how the creative use of remote laboratories can provide fun and engaging learning experiences during times of adversity.

Project-based learning

With the connectivity attribute, teachers can encourage students to participate in peer learning even though they are physically apart. Peer learning is an essential part of online learning as it promotes critical reflections through interactions among learners [13]. This is particularly important in times of emergency when face-to-face communication is hindered. Students can use collaborative platforms, such as Google Classroom and Adobe Connect, to carry out group

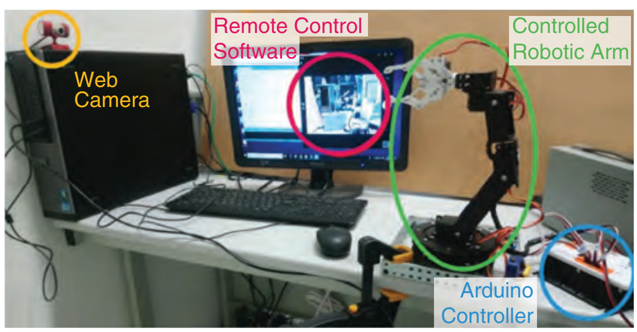

FIGURE 5. One set of remote-controllable robotic equipment developed by the Department of Mechanical and Automation Engineering at CUHK.

projects. These platforms can eliminate the limitations of space and time, allowing students to discuss and share their computer screens, documents, and images with one another. They promote connectivity among students and help learning during crises.

Upper-division and postgraduate-level courses

With the endurance attribute, teachers should encourage students to continue upper-division and postgraduate-level courses, although they are more difficult. They also need to nurture students' interest and learning motivation through online courses, which are important for lifelong learning [14]. Students can take professional online courses, which are usually conducted by professors and industry experts, and contain up-to-date information. Although these courses often require fees and are designed for people working in different industries, they are also suitable for students who want to challenge themselves. These courses enable learners to gain a deep understanding of advanced subject knowledge and apply it to solve real-world problems. With these courses, students can connect themselves to professionals and experts to continue to learn in these advanced courses, even when crises occur.

Online pedagogical experiences during the COVID-19 pandemic

To evaluate pedagogical experience and challenges for students with crisis-supported online learning, we analyzed the results of qualitative scores of CTEs collected from more than 3,000 responses. In this online survey, the observed scores were similar to the previous years' in many ways; however, students reported some challenges, such as a lack of library resources, the unavailability of IT resources, and content difficulty due to social distancing. The majority of the open-ended feedback reported online learning to be flexible, interactive, clear teaching, portable, time efficient, and convenient. A small number of users reported online learning as being hard to catch up with, baffling, and difficult to conduct hardware courses with. It is suggested that stakeholders in EE can employ several intervention plans at different levels, which include 1) organizing training workshops for the teachers and students to familiarize themselves with online learning tools, 2) providing students with resources such as home lab kits for hardware labs, 3) designing flipped-online labs, and 4) leveraging learning management systems (LMSs) to plan and manage group discussions and online instructions.

Reflections on pedagogical experiences

The pedagogical experience of CUHK during the pandemic has reminded educators of several important issues concerning teaching and learning when facing adversity:

- *The tradeoffs among different learning modes*: By and large, although the switch to online learning was abrupt, our findings show that the feedback about online learning is on par with face-to-face learning. The key lesson for educators is that even though crises may disrupt traditional teaching and learning activities, they also offer valuable opportunities to innovate current pedagogy. During the pandemic, we observed that scores of teachers have successfully sustained education through the use of online teaching tools such as micromodule and online laboratories. In CTEs, many students actually found the new learning mode flexible, efficient, and more interactive than the traditional learning mode. This indicates how crises can pave the way for the implementation of online learning initiatives, which allows us to improve existing pedagogy and enhance teaching and learning effectiveness.
- *Flexible learning during the new normal*: Teachers find online learning to be less difficult than previously thought due to the use of technology. Given the opportunities brought by crises, educators can continue to leverage micromodules to enhance flexibility in learning even when crises subside as the resources required, such as the Internet and editing software, are readily available. Because teachers and students are already familiar with micromodules, the faculty found that it was effective and quick to use micromodules for resilience pedagogy during crises. The workload is also acceptable as most micromodules can be reused with only a few amendments once developed.
- *Experiential learning is difficult*: Assessing design models for experiential learning such as hands-on laboratory work, practicums, field studies, internships, and so on is challenging but still attainable. The deployment of remotely accessible labs, recording of hands-on experiments, and delivery of lab kits to the homes of students are some remedial solutions that can help to continue laboratory courses during restricted times. According to the CTEs, one bottleneck that we encountered during the pandemic was the continuation of practical courses. Although some faculty members successfully continued laboratory courses with the use of additional hardware appliances, it should be noted that the hardware requirements for these courses are high; therefore, not all teachers were willing to invest time and money in developing these online laboratories.

Overall, rather than present an overly optimistic view, we prefer to emphasize that the sustainability of the new pedagogy depends on a number of factors, including time and resources. During times of adversity, teachers can modify the traditional pedagogy with the attributes of CRP to enhance student learning so that students can excel in both theoretical and practical courses. Seizing the opportunities inherent in crises, teachers can motivate students to engage in personalized, peer, and experiential learning as well as develop students' appreciation for lifelong learning. Nonetheless, as the mode of delivery has changed greatly, teachers need to assess student learning outcomes to further improve and refine their pedagogy, which we further discuss in the next section.

Education technology and online assessments

Assessments are important for providing accountability to the various stakeholders that learners have obtained the proposed knowledge and skills. We investigated the use of education technology for online assessment in times of crisis. Although supervised assessments are often considered more secure, fair, and effective, they are not an option when social distancing and facility closures are in effect. Given the current state of the pandemic, having a fair and objective online assessment is vital in measuring the effectiveness of teaching performance and learning outcomes. To tackle academic misconduct, online proctoring systems are a key and integral component for the constructive evaluation of knowledge acquired by learners during crises. Based on the features and technology implemented, we compared 12 popular proctoring systems and summarized them for stakeholders selecting the most appropriate one for their situation.

Online assessments

Online assessments of remote learning are a novel experience and will require great effort and technology support to replace traditional methods. The key challenges include limited physical presence, changing the production and delivery of assessment materials from teachers, the authentication of test takers for the remote assessment, and so forth. Although there are various ways to conduct online assessments, universities focus mainly on two primary methods: online formative and online summative assessments [15]. Both types of assessment can be integrated into most of the online proctoring systems to provide stakeholders the advantages of safe and effective accessibility. In the following sections, we compare and summarize the most commonly used online proctoring systems

Online proctoring systems

Online proctoring systems based on characteristics such as target user, LMS integration, the equipment needed, and so on are reported in Table 1. These systems can be further classified into three main types: 1) automated proctoring solutions, which can authenticate test takers using combined technology during different stages (e.g., preassessment and ongoing assessments and postassessment) of the online assessment [16]; 2) browser lockdown systems, which prevent test takers from entering other web pages to access information during the examination; and 3) live proctoring solutions, which involve the intervention of a human proctor during the exam or postexam. The following sections provide an overview and evaluation of these systems.

Evaluation of online proctoring systems

We conducted a review of proctoring systems using a qualitative approach and collected data from the following three sources: 1) typical proctoring system descriptions on the web, 2) published review articles and online review reports, and 3) an evaluation of some proctoring systems that CUHK has obtained for pilot testing. The cells with "N/A" indicate information that we were unable to obtain. Although we collected information on proctoring systems from these three sources, we also developed a basic evaluation framework for educational institutions to vet the proctor systems [17].

We used seven attributes, i.e., technology, cost, accuracy, error incidence, device required, ease of use, and social acceptability, in our basic evaluation framework to select proctoring systems (see Table 2). We combined these attributes and made every effort to standardize the evaluation. Due to limited time and manpower during the initial phase of the pandemic, we focused only on those systems we deemed most suitable and appropriate to our situation. Nonetheless, it is worth mentioning that we also actively engaged with our sister institutions through the community of practice-sharing events in learning more about these systems.

CUHK has developed an evaluation framework during the process of identifying and selecting the appropriate online assessment proctoring system during crises. The following are four basic functions, listed in order of priority, that we used as selection criteria:

1) *Security*: The protection of students' records and privacy is the utmost priority during the vetting process of multiple proctoring systems. CUHK has a specific and rigorous data policy that emphasizes the importance of keeping students' information private and secure. It also requires end-to-end security and that all data encrypted during the whole process, which includes pre-assessment, ongoing assessment, and postassessment.
2) *Compatibility*: The compatibility of the LMS, which needs to be consistent with universities' systems, is vital during the vetting process among multiple proctoring systems. Although proctoring systems claim to be compatible with most LMSs, it is critical for universities to test the systems themselves using a gradual increase in capacity to ensure a smooth ramp up.
3) *Ease of use*: This factor assesses the ease of use for the various stakeholders, the minimum setup operation for online assessments, and the effort required for learners to attend remote exams. It contributes to setting up online assessments efficiently, especially for those with nontechnical skills. The selected 12 proctoring systems (see Table 2) are all manageable and effective for various stakeholders to set up and join online assessments.
4) *Robustness*: This function appraises the system's ability to cope with failures and breakdowns. It also assesses whether the systems could operate stably when accommodating a surging volume of test takers in a short period. To implement robust online assessment, CUHK uses Respondus Monitor to ensure sustained operation in peak time.

Reflections on online proctoring systems

The selection and use of online proctoring systems have raised many concerns:

- *Seamless integration is still elusive*: Although most of the proctoring systems work independently, many still do not

Table 1. The 12 online proctoring systems.

System	URL	Target User	LMS Integration	Equipment Needed	Feature Summary	Types
ProctorU	https://bit.ly/3eCB0df	Professional organizations and academic institutions, e.g., Cameron University	Blackboard, Canvas, Brightspace by D2L, and Moodle	Webcam	Automatic proctoring service provider and a live proctor	Automated and live proctoring solutions
Examity	https://bit.ly/2VfbZxd	Companies and higher-education institutions, e.g., Indiana University	Blackboard, Canvas, D2L, Moodle, and Sakai	Webcam	Artificial intelligence-(AI)-powered and integrated proctoring solutions	Automated proctoring solutions
ExamSoft	https://bit.ly/2AfeUP4	Corporate environments, governments, and higher-education institutions, e.g., the University of Rhode Island College of Pharmacy	Blackboard, Canvas, and integration with other LMSs	Webcam	AI-powered and human-reviewed proctoring solutions	Automated and live proctoring solutions
ProctorFree	https://bit.ly/2BCs0pY	Education institutions and corporate environments, e.g., Gracor Languages Services	Blackboard, Canvas, Brightspace by D2L, and Moodle	Webcam	Automatic proctoring service provider	Automated proctoring solutions
Kryterion	https://bit.ly/3dyDJmH	Education institutions and companies	Blackboard, Moodle, and integration with other LMSs	Webcam	Global testing solutions and live proctors	Automated proctoring solutions, Browser lockdown systems, and live proctoring solutions
Proctorio	https://bit.ly/3eByLqK	Corporate environments, governments, and education institutions, e.g., the University of California, Berkeley	Blackboard, Canvas, Moodle, and Brightspace by D2L	Webcam	A comprehensive learning-integrity platform	Automated and live proctoring solutions
ProctorEdu	https://bit.ly/3ezDzvZ	Corporate environments and education institutions, e.g., Financial University, Sberbank	YouTestMe, Moodle, MindScroll, and custom integration	Webcam	AI-powered and human-reviewed proctoring solutions	Automated and live proctoring solutions
AI Proctor	https://bit.ly/3ie1qo2	Education institutions	Canvas, Moodle, and integration with other LMSs	Webcam with a minimum 640 × 480 resolution	Technologically advanced online proctoring solution and live proctors	Automated and live proctoring solutions
Proctortrack	https://bit.ly/2CJ6fFz	Education institutions, corporate certification, and human resources pre-employment, e.g., Blackboard, Sakai, and edX	Blackboard, integration with linear time-invariant- and application programming interface-based platforms, Moodle, and integration with other LMSs	Webcam with minimum 800 × 600 resolution	Live proctoring using AI	Automated and live proctoring solutions
RecordTS	https://bit.ly/31he0wv	Corporate environments and education institutions, , e.g., the Massachusetts Institute of Technology	N/A	N/A	A Windows remote desktop session recording software that can monitor user-activity premises or in the cloud	Browser lockdown systems
Pearson VUE	https://bit.ly/2BI9Hj2	Test owners and test takers in all industries, e.g., academic and admissions and financial services	N/A	Classroom, computer lab, mobile testing, and public test center	Computer-based testing for programs	Browser lockdown systems
Gradescope	https://bit.ly/36v5iwG	Education institutions	Blackboard, Canvas, Brightspace, Sakai, and Moodle	N/A	AI-assisted grading, grade-written exams, grade homework, and run code auto graders	Browser lockdown systems

N/A: not applicable.

integrate seamlessly with LMSs. Consequently, customizable solutions with specialized integration had to be performed, resulting in additional work for technical staff and teachers to ensure a smooth operation.

- *Robust data accessibility*: Due to network connectivity, which varies in different parts of the world, some proctoring systems may have erratic online availability due to a lack of data centers or cloud computing capabilities. To ensure data availability, it is recommended that universities prioritize proctoring systems that connect to a robust global ecosystem with a scalable infrastructure.
- *Identifying tailor-made options*: A good combination of options can provide educational institutions with flexible, automated, and cost-effective proctoring solutions. Stakeholders can refer to the basic evaluation framework for assistance with selecting the appropriate proctoring system based on their volume of online assessments, the budget, and other requirements [17].

Table 2. A comparison of technologies in online proctoring systems.

Category	Technology	Cost	Accuracy	Error Incidence	Device Required	Ease of Use	Social Acceptability	Proctoring System
Biological biometrics	Face recognition	Medium	High	Pose and lighting variations	Camera	High	Medium	ProctorU, Examity, ExamSoft, ProctorFree, Kryterion, Proctorio, AI Proctor, and Proctortrack
	Knuckle scan	Low	Medium	Capacity of knuckle databases	Camera	High	High	Proctortrack
	Fingerprint recognition	Low	High	Scanner	Fingerprint scanner	High	High	N/A
	Iris scan	High	High	Lighting	Camera	Medium	Low	N/A
	Retina scan	High	High	Diseases, e.g., cataracts	Camera	Low	Low	N/A
Behavioral biometrics	Keystroke analytics	Low	Medium	Typing speed	Keyboard	High	High	Examity and Kryterion
	Voice recognition	Low	Medium	Background noise	Microphone and telephone	High	Medium	ProctorU, ExamSoft, Kryterion, Proctorio, AI Proctor, and Proctortrack
	Gesture	Low	Medium	N/A	Camera	High	High	ProctorU and AI Proctor
	Signature	Low	Medium	Imitation	Optic pen and touch panel	High	High	N/A
Technology-enhanced approach	Lockdown browser	Low	Low	N/A	Computer	High	High	ProctorU, Examity, ExamSoft, Kryterion, Proctorio, AI Proctor, Proctortrack, RecordTS, and Pearson VUE
	Exam analytics dashboard	Medium	High	N/A	Computer	High	High	ProctorU, ExamSoft, ProctorFree, Proctorio, ProctorEdu, and Proctortrack
	Robust reporting	Low	High	N/A	Computer	Medium	High	ProctorU, ExamSoft, ProctorFree, Kryterion, Proctorio, ProctorEdu, and Proctortrack
	Remote desktop session recording	Low	Medium	Lighting	Camera	High	Medium	ExamSoft, ProctorFree, ProctorEdu, AI Proctor, Proctortrack, and RecordTS
	Video/record of the exam	Low	High	Lighting	Camera	High	Medium	ProctorU, Examity, ExamSoft, ProctorFree, Kryterion, Proctorio, ProctorEdu, AI Proctor, Proctortrack, and RecordTS
Others	Customized launch process	High	High	N/A	N/A	Low	Medium	ProctorU, Examity, Kryterion, Proctorio, ProctorEdu, AI Proctor, Proctortrack, RecordTS, and Pearson VUE
	Live proctoring	High	Medium	Qualification	Camera	Low	High	ProctorU, ExamSoft, Kryterion, Proctorio, ProctorEdu, AI Proctor, and Proctortrack

N/A: not applicable.

Conclusions

Crises in human history have come in different forms, leading to an accumulation of knowledge on crisis management in the form of handling and learning from these difficult events. Technological innovations, preparedness, and active response plans are important, albeit challenging, for organizations and stakeholders in times of crisis. During such events, the continuation of student education is the prime objective for all disciplines in the field of EE.

We outlined the CMF in engineering curriculum, which can be utilized to craft and design a systematic plan that determines policies for the planning, delivery of content, and assessment of learning outcomes during crises. We propose a CRP, which incorporates resilience attributes to sustain teaching and learning in various courses. We further examine the role of educational technologies for fair and objective learning assessment using different online proctoring systems. To improve online engineering courses during the COVID-19 pandemic, we administered a survey, identified challenges, and suggested interventions by reflecting on our experiences. The integration of the aforementioned components can deliver competent learning standards in EE that meet different learning objectives, all while providing educators, students, and other stakeholders with effective online teaching and learning alternatives during difficult situations. In times of crisis, great challenges present great opportunities that often lead to great outcomes for the "new normal."

Acknowledgments

Irwin King is the corresponding author of this article.

Authors

Irwin King (king@cse.cuhk.edu.hk) received his Ph.D. degree in computer science from the University of Southern California, Los Angeles, California. He is a professor and the chair of the Department of Computer Science and Engineering, the Chinese University of Hong Kong, Hong Kong SAR, China. He received both the ACM CIKM2019 Test of Time Award and the ACM SIGIR 2020 Test of Time Award for his work in social computing. He is an associate editor of *Journal of Neural Networks*. His research interests include machine learning, social computing, and big data. He is a distinguished member of ACM and a Fellow of IEEE.

Chandni Saxena (csaxena@cse.cuhk.edu.hk) received her Ph.D. degree in computer engineering from Jamia Millia Islamia (a central university), New Delhi, India, in 2019. Currently, she is a research assistant in the Faculty of Engineering, the Chinese University of Hong Kong, Hong Kong SAR, China, where she is with the curriculum design team of the AI4Future Project, which aims to design and develop a pretertiary artificial intelligence curriculum for the city of Hong Kong. Her research interests include network analysis, social computing, and machine learning. She is a Member of IEEE.

Christian Pak (ypak@cse.cuhk.edu.hk) received his M.A. degree in Chinese from the City University of Hong Kong, Hong Kong, China. He is a project coordinator at the Knowledge and Education Exchange Platform (KEEP) and a current graduate student at the School of Chinese, the University of Hong Kong, Hong Kong SAR, China. Prior to joining the KEEP team, he served at the School of Languages, Linguistics, Literatures, and Cultures at the University of Calgary in 2019. He is a member of the Modern Language Association of America. His research interests lie in medieval Chinese literature, alongside Chinese (particularly Middle Chinese) and Cantonese linguistics. Recently, he has been interested in exploring new pedagogical approaches using technology in language acquisition.

Chak-ming Lam (paullam0702@gmail.com) received his B.A. in English and Chinese from the University of Hong Kong, Hong Kong SAR, China. He is currently a teacher at the Delia Memorial School (Glee Path), Hong Kong SAR, China. Prior to his current position, he worked as a technical writer in the Department of Computer Science and Engineering, the Chinese University of Hong Kong, Hong Kong SAR, China, in 2020, where he engaged in technology and education writing. He is interested in exploring how technology can be utilized to enhance student learning.

Haiyan Cai (cahihaiyanstella@gmail.com) received her master's degree from the University of Hong Kong, Hong Kong SAR, China, where she is a first-year Ph.D. student. She worked as a research assistant in the Department of Computer Science and Engineering of the the Chinese University of Hong Kong, Hong Kong SAR, China from 2019 to 2020. Her research interests include education technology and computational thinking.

References

[1] A. Boin, "The new world of crises and crisis management: Implications for policymaking and research," *Rev. Pol. Res.*, vol. 26, no. 4, pp. 367–377, 2009. doi: 10.1111/j.1541-1338.2009.00389.x.

[2] R. Winthrop, "COVID-19 and school closures: What can countries learn from past emergencies?" Brookings, Washington, D.C., 2020. [Online]. Available: https://www.brookings.edu/research/covid-19-and-school-closures-what-can-countries-learn-from-past-emergencies/

[3] L. Buchanan, Q. Bui, and J. K. Patel. "Black Lives Matter May Be the Largest Movement in U.S. History." *NY Times,* July 2020. https://www.nytimes.com/interactive/2020/07/03/us/george-floyd-protests-crowd-size.html

[4] "UNSDG—Policy brief: Education during COVID-19 and beyond," United Nations Sustainable Development Group, Aug. 2020. [Online]. Available: https://unsdg.un.org/resources/policy-brief-education-during-covid-19-and-beyond

[5] J. T. Mitchell, "Crisis intervention & CISM: A research summary," vol. 12, p. 2004, Aug. 2003. [Online]. Available: http://www.icisf.org/articles/cism_research_summary.pdf

[6] D. Smith, "Beyond contingency planning: Towards a model of crisis management," *Ind. Crisis Quart.*, vol. 4, no. 4, pp. 263–275, 1990. doi: 10.1177/108602669000400402.

[7] J. Salter, "Towards a better disaster management methodology," *Aust. J. Emerg. Manage.*, vol. 10, no. 4, pp. 8–16, 1995.

[8] M. Ungar, "Systemic resilience: Principles and processes for a science of change in contexts of adversity," *Ecol. Soc.*, vol. 23, no. 4, p. 34, 2018. doi: 10.5751/ES-10385-230434.

[9] R. Chow, C. M. Lam, and I. King, "Crisis resilience pedagogy (CRP) for teaching and learning," in *Proc. IEEE Int. Conf. Eng., Technol. Educ. (TALE*), to be published.

[10] H. Alamri, V. Lowell, W. Watson, and S. L. Watson, "Using personalized learning as an instructional approach to motivate learners in online higher education: Learner self-determination and intrinsic motivation," *J. Res. Technol. Educ.*, vol. 52, no. 3, pp. 322–352, 2020. doi: 10.1080/15391523.2020.1728449.

[11] N. P. Rakhashia, A. A. Bhurane, and V. M. Gadre, "Teaching signals and systems - a first course in signal processing," in *Proc. 2020 IEEE Int. Conf. Acoustics, Speech Signal Process. (ICASSP)*, pp. 9224–9228. doi: 10.1109/ICASSP40776.2020.9054231.

[12] D. Han and M. Leung, "Flipped online laboratory for making students' first robot," The Chinese Univ. of Hong Kong, 2020, unpublished.

[13] Z. Alt inay, "Evaluating peer learning and assessment in online collaborative learning environments," *Behav. Inf. Technol.*, vol. 36, no. 3, pp. 312–320, 2017. doi: 10.1080/0144929X.2016.1232752.

[14] M. Peters and M. Romero, "Lifelong learning ecologies in online higher education: Students' engagement in the continuum between formal and informal learning," *Brit. J. Educ. Technol.*, vol. 50, no. 4, pp. 1729–1743, 2019. doi: 10.1111/bjet.12803.

[15] K. A. Gamage, E. K. de Silva, and N. Gunawardhana, "Online delivery and assessment during COVID-19: Safeguarding academic integrity," *Educ. Sci.*, vol. 10, no. 11, p. 301, 2020. doi: 10.3390/educsci10110301.

[16] Y. Atoum, L. Chen, A. X. Liu, S. D. Hsu, and X. Liu, "Automated online exam proctoring," *IEEE Trans. Multimedia*, vol. 19, no. 7, pp. 1609–1624, 2017. doi: 10.1109/TMM.2017.2656064.

[17] H. Cai and I. King, "Education technology for online learning in times of crisis," in *Proc. IEEE Int. Conf. Eng., Technol. Educ. (TALE)*, to be published.

SP

Empowering remote presence for industry and society

IEEE MMSP 2021 will be held on October 06–08, 2021 in Tampere, Finland. It is 23rd in the series, organized by the Multimedia Signal Processing Technical Committee of the IEEE Signal Processing Society (SPS), with the aim to bring together researchers and practitioners from academia and industry, passionate about multimedia signal processing, to share their knowledge, exchange ideas, explore future research directions and network.

Prospective authors are invited to submit regular and special session papers (full length, 4-6 pages) and demo papers (1 page). Conference content will be submitted for inclusion into IEEE Xplore as well as other Abstracting and Indexing (A&I) databases. Papers are solicited in, but not limited to, the following areas:

- Image / video compression
- Image / video restoration and processing
- Machine listening for analysis of acoustic scenes
- Speech / music / audio processing
- Sound source localization and separation
- Multimedia communications and networking
- Multimedia quality assessment
- Internet of Things (IoT)-based multimedia systems and applications
- Multimedia hardware design
- Multimedia databases and digital libraries
- Multimedia security and forensics
- Multimedia interaction
- Multimedia big data analytics
- Deep Learning for Multimedia
- Multimedia processing for tele-medicine and health care
- Augmented, mixed and virtual reality
- Multimedia systems for emerging applications

Timeline

- Special Session / Grand Challenge proposal submission: 30 April 2021
- Regular and Special Session paper submission: 28 May 2021
- Notification of paper acceptance: 16 July 2021
- Demo paper submission: 13 August 2021

General Chairs
Atanas Gotchev, *Tampere University, Finland*
Dong Tian, *InterDigital, USA*

Technical Programme Chairs
Annamaria Mesaros, *Tampere University, Finland*
Mårten Sjöström, *Mid Sweden University, Sweden*
Gene Cheung, *York University, Canada*
Giacomo Boracchi, *Politecnico di Milano, Italy*
Dan Zeng, *Shanghai University, China*

Keynote Chair
Christine Guillemot, *INRIA, France*

Special Sessions Chairs
Federica Battisti, *University Roma3, Italy*
Patrick Le Callet, *University of Nantes, France*

Demo / Industrial Chair
Sebastian Schwarz, *Nokia Technologies, Germany*

Grand Challenges Chair
Jiaying Liu, *Peking University, China*

Publicity Chair
Pasi Pertilä, *Tampere University, Finland*

Finance Chair
Sari Peltonen, *Tampere University, Finland*

Publication Chair
Joao Ascenso, *University of Lisbon, Portugal*

International Liaison Chair
Ricardo de Queiroz, *University of Brazil, Brazil*

Local Arrangements Chairs
Sarianne Niemelä, *Tampere University, Finland*

Awards Chairs
Jörn Ostermann, *University of Hannover, Germany*
Thomas Maugey, *INRIA, France*

Advisory Committee Chair
Antonio Ortega, *University of Southern California, USA*

IEEE Signal Processing Society

www.mmsp2021.org

Cover photo by Laura Vanzo, Visit Tampere

Digital Object Identifier 10.1109/MSP.2021.3070730

Çağatay Candan

Proper Definition and Handling of Dirac Delta Functions

Dirac delta functions are introduced to students of signal processing in their sophomore year. Quite understandably, Dirac delta functions, which should be more aptly called *generalized functions* or *distributions*, cannot be comprehensively given to a young audience at the beginning of their engineering education. Instead, a simplified and abridged definition is presented, and the implications of the definition in signal processing problems are illustrated through numerous examples, following the footsteps of Oppenheim et al. [1], [2].

Students typically learn the properties by developing an affinity through their usage. As their mathematical knowledge matures, some students tend to notice inconsistencies related to the sugarcoated definitions and start questioning the mathematics behind them. Unfortunately, the inquisitive questions of these students are rather difficult to answer convincingly due to the lack of sources on generalized functions at the level of undergraduate/graduate engineering students. The goal of these notes is to scratch the sugarcoating a bit and provide the basics of generalized functions, limits, and derivatives as well as their usage in signal processing problems.

As an illustrative example, the Fourier transform of $f(t) = 1$, which is $F(\Omega) = 2\pi\delta(\Omega)$, is typically "proven" with the application of the inverse Fourier transform on $F(\Omega) = 2\pi\delta(\Omega)$. However, according to the standard calculus results, the Fourier transform of $f(t) = 1$, which is $\mathcal{F}\{1\} = \int_{-\infty}^{\infty} \exp(-j\Omega t)\,dt$, ceases to exist for any Ω in the ordinary calculus sense. The plot further thickens when the Fourier transform of the unit step function, sign function, and Hilbert transform discussions come into play.

Generalized functions enable these calculations, and they are indispensable tools of our field, yet their proper understanding, true definition, and the whys and hows about their usage require an update to our classical calculus knowledge. Such an update, however incomplete, is the topic of this lecture note.

Relevance

Paul Dirac is one of the giants among the great physicists of the early 20th century. It is a compliment to our profession that he received his first academic degree in electrical engineering (from the University of Bristol). He said,

> I owe a lot to my engineering training because it [taught] me to tolerate approximations. Previously to that I thought . . . one should just concentrate on exact equations all the time. Then I got the idea that, in the actual world, all our equations are only approximate. We must just tend to greater and greater accuracy. In spite of the equations being approximate, they can be beautiful

The function $\delta(t)$ introduced by Dirac is now called the *Dirac delta function*; it provides great computational and conceptual advantages in calculations involving diverging integrals, which is the case for some Fourier integrals. In addition, the inclusion of the Dirac delta function to the calculus of ordinary functions enables the differentiation of discontinuous (generalized) functions, paving the way toward a consistent analysis of highly practical engineering problems, such as circuit theory problems involving switches, unified treatment for mixed random variables (random variables that are both discrete and continuous), and more.

Despite the abundance of topics utilizing Dirac delta functions in signal processing, there are only a few sources explaining the true nature of the approximation involved to the signal processing audience [3, Appendix I], [4]. This column is prepared to answer some of the questions on generalized functions, illustrate their properties, and show their proper usage in some signal processing calculations.

The intended audience of this lecture note includes instructors, researchers with an inclination toward theory, and graduate students getting close to fulfilling their course requirements, say, studying for Ph.D. qualification exams.

Digital Object Identifier 10.1109/MSP.2021.3055025
Date of current version: 28 April 2021

 1053-5888/21©2021IEEE

For beginners to the topic, the author suggests following the mainstream track and developing an affinity for the topic first by following the wisdom of Oppenheim et al. [1], [2].

The conventional treatment aims to develop a working knowledge of Dirac delta functions, which is a noteworthy goal on its own, and gives a good "first-order approximation" to the topic. Science and engineering are built upon successively refined approximations, which Paul Dirac has alluded to as a potential source of beauty.Especially in engineering, approximate models/explanations are important, beyond their aesthetic value, because of the basic need for working tools and methods for the solution of practical problems.

In a typical undergraduate course, the need for a working solution may easily overshadow the need for a comprehensive theoretical treatment. As an example, the first course in physics studies the mechanics of inclined planes, stacked boxes with high/low friction surfaces, and so on. If we consider two stacked wood blocks on a flat surface, we may say that the weight of top block is balanced with the normal force so that the net force on the block is zero. This comment can be used to explain why two blocks do not coalesce into a single piece.

However, if we think about the nature of the normal force, it is typically explained as a direct consequence of Newton's laws of motion (the law of action–reaction), and Newton's laws are brought upon students axiomatically in relation to Newton's empirical observations. Hence, the contents of Physics 101 correctly predict that two stacked wood blocks will not coalesce into a single piece without saying much about the mechanism behind the process!

In spite of that, Physics 101 students learn to use and appreciate the benefit of defining a normal force through a series exercises and problems, just like a beginner signal processing student working his or her way through a set of exercise problems on Dirac delta functions. Much later, physicists with advanced degrees learn that the macroscopic normal force is due to the Pauli exclusion principle applied to bulk matter [5]. Needless to say, such a comprehensive answer is of no help to Physics 101 students working on problems with inclined planes.

The situation is almost analogous for signal processing students and Dirac delta functions. Hence, the author believes that exposure to Dirac delta functions beyond the conventional Oppenheim et al. level can be safely postponed to graduate studies. Of course, professionals in the field, lecturers, and researchers can refer quick learners with inquisitive questions to this lecture note, disregarding the suggested timeline.

Prerequisites

The only prerequisites are a working knowledge of freshman calculus, basic signal processing theory, and a keen eye for detail.

Problem statement

The main focus is on the handling of integrals, limits, and derivatives that do not exist in the standard calculus sense. The Fourier transform of $u(t)$ (the unit step function), $F(\Omega) = \int_{-\infty}^{\infty} u(t)\exp(-j\Omega t)\,dt$, is the prime illustrative example. This Fourier transform integral requires the evaluation of $\int_0^{\infty} \cos(\Omega t)\,dt$ and $\int_0^{\infty} \sin(\Omega t)\,dt$, which are known to diverge according to standard calculus results. However, signal processing textbooks express the result as $\mathcal{F}\{u(t)\} = 1/j\Omega + \pi\delta(\Omega)$ [1, Table 4.2].

The appearance of $\delta(\cdot)$ function hints at the divergence of the Fourier integral to an experienced eye, but this is not the case for all divergent integrals. The Fourier transform of $\mathrm{sgn}(t)$ (the sign function) requires the evaluation of $\int_0^{\infty} \sin(\Omega t)\,dt$, which is a divergent integral. However, textbooks state that $\mathcal{F}\{\mathrm{sgn}(t)\} = 2/j\Omega$. The main problem is that the transform pair for both functions is not valid in the ordinary calculus sense but valid in the generalized sense or in the sense of distributions. This article studies the definition of generalized functions and their use in signal processing problems.

Solution

We first present some basic definitions to better explain the upcoming definitions of the Dirac delta and other generalized functions.

Function

Functions, as defined on the set of real numbers, map real numbers to real numbers. Functions are interpreted in a pointwise manner. For example, $\phi(t) = t^2$ maps t_0 in $(-\infty,\infty)$ to t_0^2 in $[0,\infty)$.

Linear functional

A functional is a mapping from the space of functions to real numbers. For example, the area functional defined as $\mathrm{Area}\{\phi\} = \int_{-\infty}^{\infty} \phi(t)\,dt$ maps the function $\phi(t)$ to the numerical value of the total area under $\phi(t)$. A functional that satisfies the linearity conditions (homogeneity and additivity, [1, Sec. 1.6.6]) is called a *linear functional*. Our focus is entirely on linear functionals. Hence, the term *functional* should be interpreted as a linear functional in these notes.

It is easy to verify that the functional $T_f\{\cdot\}$,

$$T_f\{\phi(t)\} \triangleq \langle f(t), \phi(t) \rangle = \int_{-\infty}^{\infty} f(t)\phi(t)\,dt, \tag{1}$$

satisfies the conditions of linearity. We use the notations $T_f\{\phi(t)\}$ and $\langle f(t), \phi(t)\rangle$ interchangeably to denote functionals. $T_f\{\phi(t)\}$ explicitly shows that the "input" $\phi(t)$ is mapped to an "output," i.e., a real number. The function $f(t)$ appearing in the subscript of $T_f\{\phi(t)\}$ characterizes the mapping. As an example, the area functional, previously given, can be realized by substituting $f(t)$ with 1 in (1). The second notation $\langle f(t), \phi(t)\rangle$ is handy in many calculations due to the symmetry between $f(t)$ and $\phi(t)$ in (1).

We refer to the function $\phi(t)$ as the test function. Hence, $T_f\{\phi(t)\}$ is said to operate on test functions. Generalized functions or distributions, shown as $f(t)$, are built upon the "observed" action of functionals on the test functions, as described in the next section.

Generalized equality

If functions $f(t)$ and $g(t)$ induce the same functional, that is, $T_f\{\phi(t)\}$ and $T_g\{\phi(t)\}$ yield identical outputs for all test functions, then functions $f(t)$ and $g(t)$ are said to be equal in the generalized sense. We show the generalized equality with the notation of $f(t) \overset{(g)}{=} g(t)$:

$$f(t) \overset{(g)}{=} g(t) \Leftrightarrow \langle f(t), \phi(t)\rangle = \langle g(t), \phi(t)\rangle \quad \text{for all } \phi(t). \tag{2}$$

To make the statements precise, we need to specify the function class for the test functions and also give a discussion of Lebesgue integration. We refer readers to [6, Ch. 6] for a readable account of these topics. As readers can intuitively appreciate, the class for the test functions should be sufficiently "rich" and "refined" so that the generalized equality in (2) presents practically useful results. For example, if the test functions are limited to constant functions, say, $\phi(t) = c$, where c is a real number, the generalized equality in (2) only implies the equality of the area under two functions, which is of rather limited value.

In this text, we assume that the test function class is infinitely differentiable functions in the form of Gaussian functions:

$$\phi_{\mu,\sigma}(t) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(t-\mu)^2}{2\sigma^2}\right), \tag{3}$$

with arbitrary mean μ and spread σ. We take this class of test functions as sufficiently rich and refined so that the generalized equality $f(t) \overset{(g)}{=} g(t)$ in (2) becomes practically meaningful. [The class of infinitely differentiable test functions with rapid decay at infinity is called the *Schwartz space* [6], [7]. The Hermite functions, an orthonormal and complete set for L_2, are members of this class. Laurent Schwartz received the Fields Medal in 1950 for building the mathematical foundation (theory of distributions) to the framework of Dirac.]

Dirac delta function

We consider a specific functional, called the *evaluation functional*, that maps the function $\phi(t)$ to $\phi(t_0)$, i.e., the evaluation functional maps $\phi(t)$ to the value of its sample at $t = t_0$. The evaluation functional is clearly linear, but it is not possible to express the evaluation functional in the form of (1) with a regular $f(t)$ function. In spite of that, we substitute $f(t)$ with $\delta(t-t_0)$ in (1) and use the following as a formal definition of the evaluation functional:

$$\int_{-\infty}^{\infty} \delta(t-t_0)\phi(t)\,dt = \phi(t_0) \quad \text{for all } \phi(t). \tag{4}$$

We do not question the existence of the $\delta(t)$ function at this point but treat it as a regular function for now. Readers may interpret (4) as another notation for the evaluation functional from which some properties, such as the linearity of the functional, can be readily observed. Our goal is to derive some properties of $\delta(t)$, given in Table 1, first and then answer existence questions.

Verification of the multiplication property

Let's study the product of $f(t)$ and $\delta(t-t_0)$, which is $f(t_0)\delta(t-t_0)$ according to the multiplication property $f(t)\delta(t-t_0) \overset{(g)}{=} f(t_0)\delta(t-t_0)$ in Table 1. To prove the generalized inequality, we need to show that $\langle f(t)\delta(t-t_0), \phi(t)\rangle = \langle f(t_0)\delta(t-t_0), \phi(t)\rangle$ for all test functions. We focus on the term on left-hand side, $\langle f(t)\delta(t-t_0), \phi(t)\rangle$, first:

$$\begin{aligned}\langle f(t)\delta(t-t_0), \phi(t)\rangle &= \int_{-\infty}^{\infty} f(t)\delta(t-t_0)\,\phi(t)\,dt\\ &\overset{(a)}{=} \int_{-\infty}^{\infty} \delta(t-t_0)\,\hat{\phi}(t)\,dt\,\Big|_{\hat{\phi}(t)=f(t)\phi(t)}\\ &\overset{(b)}{=} \hat{\phi}(t_0)\\ &= f(t_0)\phi(t_0).\end{aligned} \tag{5}$$

In line (a), $\hat{\phi}(t) = f(t)\phi(t)$ is introduced, and $\hat{\phi}(t)$ is assumed to be a member of the test function class due to its "richness" and "fineness." Line (b) is due to the definition of evaluation functional.

The right side of equality $f(t)\delta(t-t_0) \overset{(g)}{=} f(t_0)\delta(t-t_0)$ can be worked out as follows:

$$\begin{aligned}&\langle f(t_0)\delta(t-t_0), \phi(t)\rangle\\ &\quad= \int_{-\infty}^{\infty} f(t_0)\delta(t-t_0)\phi(t)\,dt\\ &\quad= f(t_0)\int_{-\infty}^{\infty} \delta(t-t_0)\phi(t)\,dt\\ &\quad= f(t_0)\phi(t_0).\end{aligned} \tag{6}$$

Combining (5) and (6), we have

$$\langle f(t)\delta(t-t_0), \phi(t)\rangle = \langle f(t_0)\delta(t-t_0), \phi(t)\rangle, \quad \text{for all } \phi(t), \tag{7}$$

which concludes the proof of $f(t)\delta(t-t_0) \overset{(g)}{=} f(t_0)\delta(t-t_0)$.

An important takeaway message from the proof of the first property is not the final result but the proof procedure followed for the generalized

Table 1. The properties for the Dirac delta function and its derivatives.

Basic	Multiplication	$f(t)\delta(t-t_0) \overset{(g)}{=} f(t_0)\delta(t-t_0)$
	Scaling	$\delta(at) \overset{(g)}{=} \frac{1}{\lvert a\rvert}\delta(t)$
	Sifting	$\int_{-\infty}^{\infty} f(t)\delta(t-t_0)\,dt = f(t_0)$
	Convolution	$\delta(t) * f(t) \overset{(g)}{=} f(t)$
Advanced	Multiplication	$f(t)\delta^{(n)}(t-t_0) \overset{(g)}{=} \sum_{k=0}^{n}(-1)^k \binom{n}{k} f^{(k)}(t_0)\delta^{(n-k)}(t-t_0)$, where $\frac{d^n}{dt^n}\delta(t) = \delta^{(n)}(t)$, and $\frac{d^n}{dt^n}f(t) = f^{(n)}(t)$
	Scaling	$\delta(f(t)) \overset{(g)}{=} \sum_{k=1}^{K} \frac{1}{\lvert f'(t_k)\rvert}\delta(t-t_k)$, where t_k are zeros of $f(t)$, i.e., $f(t_k) = 0, k = \{1, 2, \ldots, K\}$
	Sifting	$\int_{-\infty}^{\infty} f(t)\delta^{(n)}(t-t_0)\,dt = (-1)^n f^{(n)}(t_0)$
	Convolution	$f(t) * \delta^{(n)}(t) \overset{(g)}{=} f^{(n)}(t)$

equality. The equality sign $\overset{(g)}{=}$ appearing in $f(t) \overset{(g)}{=} g(t)$ denotes the equality of the functionals for every member of the test function class. It is, indeed, very different from the ordinary equality sign.

A rather silly, but memorable, analogy given by one of my instructors can be repeated as follows: Assume that you are in a county fair, and there is a contest to identify an unknown animal. Contestants are allowed to ask only yes/no questions. After several rounds of questions, you learn that the animal is green, lives in a lake, is capable of leaping significant distances, and quacks. Given this information, can you say that the animal is a frog?

If you have asked a sufficiently large number of informative questions (the richness and fineness of the question class), you can be pretty sure that the animal is a frog! However, there is always a possibility that the animal is of another species that is capable of imitating a frog quite closely! If you are only interested in the actions of this animal, though, there is no harm in calling the animal, irrespective of its genus, a frog or a generalized frog!

Analogous to the story, a generalized function $f(t)$ is characterized by its response to the probing test functions $\phi(t)$. Generalized functions are declared equal if they give the same response to all test functions.

The major mishap in the treatment of the impulse function or Dirac delta function in all signal processing texts is the usage of an ordinary equality sign $=$ instead of a generalized equality sign $\overset{(g)}{=}$. This carries the potential of interpreting equations involving $\delta(t)$ in a pointwise manner, which is prone to inconsistencies and calculation mistakes.

Verification of the scaling property

Let's verify the scaling property $\delta(at) \overset{(g)}{=} (1/|a|)\delta(t)$, given in Table 1. The left side of the equality can be written as

$$\begin{aligned}\langle \delta(at), \phi(t)\rangle &= \\ &\int_{-\infty}^{\infty} \delta(at)\phi(t)\,dt\,\Big|_{u=at} \\ &= \frac{1}{|a|}\int_{-\infty}^{\infty}\delta(u)\phi\Big(\frac{u}{a}\Big)du = \frac{\phi(0)}{|a|}. \end{aligned} \tag{8}$$

Here, $\phi(u/a)$ is assumed to be in the test function class, as in the proof of the first property, and we have treated $\delta(at)$ as a regular function and changed the integration variable from t to $u = at$ without any due diligence (more on this later).

The right side of the equality $\delta(at) \overset{(g)}{=} (1/|a|)\delta(t)$ can be written as

$$\begin{aligned}\Big\langle \frac{1}{|a|}\delta(t), \phi(t)\Big\rangle &= \frac{1}{|a|}\langle \delta(t), \phi(t)\rangle \\ &= \frac{\phi(0)}{|a|}. \end{aligned} \tag{9}$$

Equations (8) and (9) imply the generalized equality of $\delta(at) \overset{(g)}{=} (1/|a|)\delta(t)$. Note that setting $a = -1$ in the scaling property gives $\delta(t) \overset{(g)}{=} \delta(-t)$, which is the evenness of function $\delta(t)$ in the generalized sense.

Generalized limit

Up to this point, we have averted the existence questions on the $\delta(t)$ function but, rather, focused on its properties. Now, we present a limit argument for the construction of the Dirac delta function. The described limit operation is called the *generalized limit*. In standard textbooks, the Dirac delta function is introduced as the pointwise limit of ordinary functions, which is not the correct definition and the root cause of confusion in many discussions.

The generalized limit of ordinary functions $f_n(t)$ is said to be a generalized function $f(t)$, if

$$\lim_{n\to\infty}\int_{-\infty}^{\infty} f_n(t)\phi(t)\,dt = \int_{-\infty}^{\infty} f(t)\phi(t)\,dt \tag{10}$$

is satisfied for all test functions $\phi(t)$. We denote the generalized limit as $f_n(t) \overset{(g)}{\to} f(t)$.

The Dirac delta function can be given as the generalized limit of ordinary $f_n(t)$ functions defined as follows:

$$f_n(t) = \begin{cases} \dfrac{n}{\epsilon}, & -\dfrac{\epsilon}{2n} < t < \dfrac{\epsilon}{2n} \\ 0, & \text{other} \end{cases}. \tag{11}$$

From Figure 1, it can be seen that $f_n(t)$ is a pulse of duration ϵ/n centered around $t = 0$. The area under $f_n(t)$ is unity for all n. With the running assumption that the test functions $\phi(t)$ are sufficiently smooth, we can expand the function into the Taylor series around $t = 0$:

$$\begin{aligned}\phi(t) = \phi(0) + \phi'(0)t + \phi^{(2)}(0)\frac{t^2}{2} \\ + \text{h.o.t.} \end{aligned} \tag{12}$$

Here, h.o.t. refers to the higher-order terms of the Taylor series expansion. As $n\to\infty$, the support of function $f_n(t)$, as shown in Figure 1, approaches zero. Hence, the product $\phi(t)f_n(t)$ can be approximated with the first term of the Taylor series expansion, which is $\phi(0)f_n(t)$, for large enough n. As a result, we have the equality of

$$\lim_{n\to\infty}\int_{-\infty}^{\infty} f_n(t)\phi(t)\,dt = \phi(0) \tag{13}$$

in the usual calculus sense. Given the generalized limit definition, this concludes the proof of $f_n(t) \overset{(g)}{\to} \delta(t)$ as $n \to \infty$.

The definition of the Dirac delta function as a generalized limit of ordinary functions is important in practice. Whenever in doubt, it is possible to replace $\delta(t)$ with the $f_n(t)$ functions in (11), solve the problem of interest, and then calculate the ordinary limit of the final result as $n\to\infty$. Readers are invited to do this calculation to have another verification of the scaling property in Table 1. Furthermore, the generalized limit definition establishes a connection with the "physical" interpretation of the Dirac delta function as a very-short-duration pulse, but readers should always keep in mind that the limit operation for getting shorter and shorter pulses is not an ordinary pointwise limit operation, as introduced in many undergraduate texts, but a generalized limit operation.

The Dirac delta definition by a generalized limit argument is not specific to $f_n(t)$ given by (11). Readers can examine [1, Problem 1.38] for some other ordinary functions for which the generalized limit is $\delta(t)$. The basic requirement is the construction of a unit area function sequence with diminishing support. It can be verified that both

$$\begin{aligned} g_n(t) &= \sqrt{\frac{n}{2\pi}}\exp(-nt^2/2) \text{ and} \\ h_n(t) &= n\,\mathrm{sinc}(nt) = \frac{\sin(\pi nt)}{\pi t} \end{aligned} \tag{14}$$

tend to $\delta(t)$ as $n \to \infty$ in the generalized sense.

Figure 2 shows the sketch of $h_n(t) = n\text{sinc}(nt)$ for different n values. The main lobe of the function $h_n(t)$ gets narrower and taller as n increases, yet, however large n is, there exist two sidelobes, with a peak value of about one-fifth the maximum value, on both sides of the main lobe. Furthermore, by fixing t to a nonzero value, say t_0, and evaluating $\lim_{n\to\infty} h_n(t_0)$, we get

$$\lim_{n\to\infty} h_n(t_0) = \frac{1}{\pi t_0} \lim_{n\to\infty} \sin(\pi n t_0),$$

which does not exist in the usual sense. Hence, $h_n(t)$ does not approach $\delta(t)$ in a manner that is as described in many undergraduate textbooks but approaches in the generalized sense or, equivalently, in the weak limit sense [8].

Generalized derivative of the Dirac delta function

The derivative of the Dirac delta function $d/dt\{\delta(t)\}$ is called the *doublet function* [1, Sec. 2.5.3]. It is no surprise that the differentiation operation in $d/dt\{\delta(t)\}$ is in the generalized sense, that is, according to the introduced generalized limit definition. To understand this operation, let's examine the response of $d/dt\{\delta(t)\}$ to a test function:

$$\int_{-\infty}^{\infty} \frac{d}{dt}\{\delta(t)\}\phi(t)dt = \delta(t)\phi(t)\Big|_{t=-\infty}^{t=\infty} - \int_{-\infty}^{\infty} \delta(t)\frac{d}{dt}\phi(t)dt = -\frac{d}{dt}\phi(t)\Big|_{t=0} = -\phi^{(1)}(0). \tag{15}$$

This calculation is based on the application of integration by parts to the leftmost side of (15). Since the test function $\phi(t)$ is a member of scaled and shifted Gaussian functions, the term $\delta(t)\phi(t)\big|_{t=-\infty}^{t=\infty}$ vanishes. The other term, the integral term of the integration-by-parts operation, can be expressed using the sifting property of the Dirac delta function. Hence, we get the defining relation for the doublet function as

$$\int_{-\infty}^{\infty} \delta^{(1)}(t)\phi(t)dt = -\phi^{(1)}(0). \tag{16}$$

Here, $\delta^{(n)}(t)$ and $\phi^{(n)}(t)$ refer to the nth derivative of the Dirac delta and test function $\phi(t)$ in the generalized and ordinary sense, respectively.

At this point, readers should be rightfully uncomfortable with the application of integration by parts with an integrand containing a Dirac delta function, as in (15). To the comfort of these readers (and also the ones still uneasy about the change of variables from t to $u = at$ in the scaling property discussion), we present an alternative proof path and suggest replacing $\delta(t)$ with the ordinary function $h_n(t)$ given in (14). The integration-by-parts operation with the substituted $h_n(t)$ function is now well defined, and the final result becomes

$$\int_{-\infty}^{\infty} \frac{d}{dt}\{h_n(t)\}\phi(t)dt = -\int_{-\infty}^{\infty} h_n(t)\frac{d}{dt}\phi(t)dt. \tag{17}$$

By taking the limit of both sides in (17) as $n \to \infty$ and using the generalized limit definition in (10), we reach the conclusion that, since $h_n(t) \overset{(g)}{\to} \delta(t)$, we have $d/dt\{h_n(t)\} \overset{(g)}{\to} \delta^{(1)}(t)$. The formal definition of $\delta^{(1)}(t)$ becomes the relation in (16).

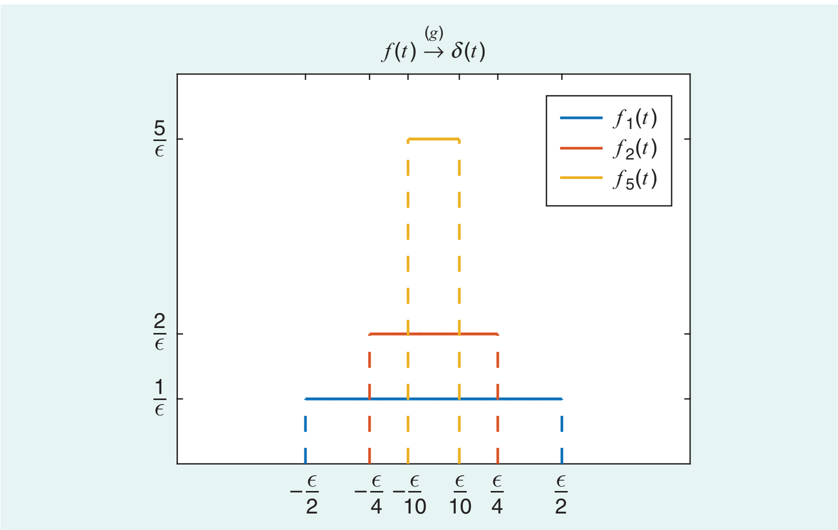

FIGURE 1. The convergence of pulse sequences $f_n(t)$ to $\delta(t)$.

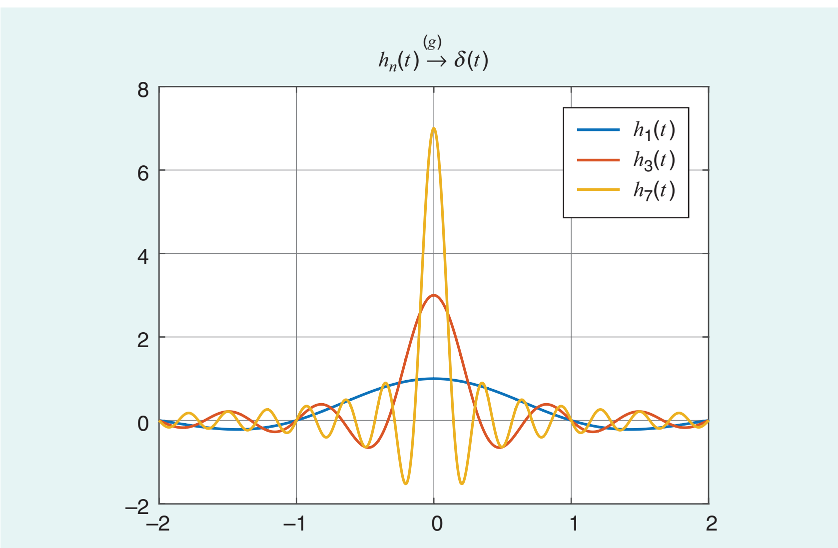

FIGURE 2. The convergence of $h_n(t) = n\text{sinc}(nt)$ to $\delta(t)$. Convergence is not in the pointwise sense!

Sifting and other properties for higher-order derivatives of the Dirac delta function are given in Table 1. These results can be called *advanced results*, since they require more than a basic understanding of the generalized functions. Many signal processing textbooks avoid these properties since even a partial justification of these results requires much more than a pictorial or pointwise justification of the $\delta(t)$ function.

Derivative of the unit step function

By replacing $h_n(t)$ with an arbitrary regular function $f(t)$ in (17), we get

$$\int_{-\infty}^{\infty} \frac{d}{dt}\{f(t)\}\,\phi(t)\,dt = -\int_{-\infty}^{\infty} f(t)\frac{d}{dt}\phi(t)\,dt. \quad (18)$$

Substituting $f(t)$ in (18) with the unit step function $u(t)$ yields

$$\begin{aligned}\int_{-\infty}^{\infty} \frac{d}{dt}\{u(t)\}\,\phi(t)\,dt &= -\int_{-\infty}^{\infty} u(t)\frac{d}{dt}\phi(t)\,dt\\ &= -\int_{0}^{\infty} \frac{d}{dt}\phi(t)\,dt\\ &\overset{(a)}{=} \phi(0)-\cancel{\phi(\infty)}\\ &= \langle \delta(t),\phi(t)\rangle,\end{aligned} \quad (19)$$

where $\phi(\infty)=0$ is used in line (a), which is due to the test function class definition. The leftmost and rightmost sides of (19) imply that $\langle (d/dt)u(t),\phi(t)\rangle = \langle \delta(t),\phi(t)\rangle$ for all test functions. This statement is equivalent to $(d/dt)u(t) \overset{(g)}{=} \delta(t)$.

From this discussion, we reach the important conclusion that an ordinary function, such as $u(t)$, when interpreted as a generalized function, has derivatives of all orders. In other words, function $u(t)$ is not a differentiable function due to its discontinuity at $t=0$, but it is differentiable for all orders in the generalized sense.

Application examples

A number of examples are presented to illustrate the application of the Dirac delta function. Our goal is to relate the applications to the generalized definitions on functions, limits, derivatives, and so on.

Example 1

Assume that a sequence $y[n]$ is formed by down-sampling $x[n]$ by two: $y[n]=x[2n]$. It is well known that the spectrum of $y[n], Y(e^{j\omega})$ is related to the spectrum of $x[n]$, $X(e^{j\omega})$, according to the relation [2, Sec. 3.6.1]

$$Y(e^{j\omega}) = \frac{1}{2}\left(X\left(e^{j\frac{\omega}{2}}\right)+X\left(e^{j\left(\frac{\omega}{2}+\pi\right)}\right)\right). \quad (20)$$

In this example, we would like to illustrate the validity of this expression for $x[n]=\exp(j\omega_0 n)$. This exercise is quite trivial from the time-domain-processing viewpoint. Since $y[n]=x[2n]=\exp(j2\omega_0 n)$, $y[n]$ is a complex exponential whose frequency is doubled after down-sampling. The frequency-domain representations of $x[n]$ and $y[n]$ are $X(e^{j\omega})=2\pi\delta(\omega-\omega_0)$, and $Y(e^{j\omega})=2\pi\delta(\omega-2\omega_0)$, respectively. This example aims to verify this basic result directly from (20).

It should be remembered that the expressions for $X(e^{j\omega})$ and $Y(e^{j\omega})$ are periodic with 2π, as the notation implies. Let's check the validity of (20) for $X(e^{j\omega})=2\pi\delta(\omega-\omega_0)$:

$$\begin{aligned}Y(e^{j\omega}) &= \frac{1}{2}\left(X\left(e^{j\frac{\omega}{2}}\right)+X\left(e^{j\left(\frac{\omega}{2}+\pi\right)}\right)\right)\\ &= \frac{2\pi}{2}\left[\delta\left(\frac{\omega}{2}-\omega_0\right)\right.\\ &\quad\left.+\delta\left(\frac{\omega}{2}+\pi-\omega_0\right)\right]\\ &= \frac{2\pi}{2}\left[\delta\left(\frac{\omega-2\omega_0}{2}\right)\right.\\ &\quad\left.+\delta\left(\frac{\omega+2\pi-2\omega_0}{2}\right)\right]\\ &\overset{(a)}{=} \frac{2\pi}{2}[2\delta(\omega-2\omega_0)\\ &\quad+2\delta(\omega+2\pi-2\omega_0)]\\ &\overset{(b)}{=} 2\pi\delta(\omega-2\omega_0).\end{aligned}$$

In line (a), we have used the scaling property of the Dirac delta function from Table 1. In line (b), the expression is rewritten to cover only a single 2π period, following the convention. As expected, the final result indeed matches the earlier result found from time-domain considerations.

Comment

The spectrum after a down-sampling operation is typically found with a frequency-domain sketch that indicates the support of $X(e^{j\omega})$ and its translated versions (see [2, Fig. 3.18]). Such a sketch is also useful to illustrate the aliasing concept. We see that when the spectrum involves a Dirac delta function, a sketch is not sufficient to explain the vanishing 1/2 coefficient in (20). We need to bring the scaling property of the Dirac delta function into play.

Example 2

Let X be a random variable with the probability density function (pdf) $f_X(x)$. The problem of interest is the pdf of the random variable $Z=X^2$.

This is a standard probability problem, and we would like to illustrate the utility of the Dirac delta function in this calculation:

$$\begin{aligned}f_Z(z) &\overset{(a)}{=} \int f_{X,Z}(x,z)\,dx\\ &\overset{(b)}{=} \int f_X(x) f_{Z|X}(z \mid X=x)\,dx\\ &\overset{(c)}{=} \int f_X(x)\,\delta(z-x^2)\,dx\\ &\overset{(d)}{=} \int f_X(x)\,\delta(x^2-z)\,dx\\ &\overset{(e)}{=} \int f_X(x)\left(\frac{\delta(x-\sqrt{z})}{2\sqrt{z}}\right.\\ &\quad\left.+\frac{\delta(x+\sqrt{z})}{2\sqrt{z}}\right)dx \quad (\text{with } z \geq 0)\\ &\overset{(f)}{=} \frac{1}{2\sqrt{z}}\left(f_X(\sqrt{z})+f_X(-\sqrt{z})\right).\end{aligned} \quad (21)$$

Line (a) is the marginalization operation. Line (b) includes a factorization for the joint density in terms of the conditional density. Line (c) introduces $Z=X^2$ into the calculation. Line (d) is due to the evenness of the Dirac delta function, $\delta(x)=\delta(-x)$. Line (e) uses the scaling property of Table 1 (from the "Advanced" section of the table). It is important to note that the integration variable in line (d) is x. Hence, for the function $\delta(x^2-z)$ appearing in the integrand, x is the variable, and z is just a constant value. Therefore, the scaling property of the Dirac delta function should be utilized by treating function x^2-z as a function of the variable x. Line (f) is due to the sifting property.

Comment

We observe that the inclusion of the Dirac delta in the operational calculus results in significant shortening of the algebra. Note that the calculation given in (21) exactly mimics a similar calculation given for the discrete random variables (probability mass functions).

More specifically, line (c) of (21) can be interpreted as follows: Let's assume that $z = 100$ and consider the integral $\int_{-\infty}^{\infty} f_X(x)\delta(z-x^2)dx$. Since the function $\delta(z-x^2)$ is equal to zero when $z \neq x^2$, this integral corresponds to checking all $x \in (-\infty, \infty)$ to find the ones satisfying the condition $x^2 = z = 100$ and "summing up" $f(x)$ values corresponding to these x values.

The main difficulty for instructors is not this interpretation but explaining the factor $1/2\sqrt{z}$, which is the Jacobian term arising during the functional mapping of random variables. The Jacobian term does not arise in discrete random variables, and the "summing up" interpretation becomes exactly correct; that is, for the probability mass functions, the sum of the probability values for x that satisfy the condition $z = x^2$ gives the probability of z. With the inclusion of the Dirac delta function in the calculus, the $1/2\sqrt{z}$ term in line (f) of (21) effortlessly comes out with the application of the scaling property.

Example 3

Let X and Y be two random variables with the joint pdf $f_{X,Y}(x,y)$. The problem is the derivation of the pdf for the random variable $Z = X + Y^2$:

$$
\begin{aligned}
f_Z(z) &\overset{(a)}{=} \int_x \int_y f_{X,Y}(x,y)\delta(z-x-y^2)dydx \\
&\overset{(b)}{=} \int_y \Big(\int_x f_{X,Y}(x,y)\delta(z-x-y^2)dx \Big) dy \\
&\overset{(c)}{=} \int_y \Big(\int_x f_{X,Y}(x,y)\delta(x-z+y^2)dx \Big) dy \\
&\overset{(d)}{=} \int_y f_{X,Y}(z-y^2, y)dy \\
&\overset{(e)}{=} \int_y f_Y(y) f_{X|Y}(z-y^2 \,|\, Y=y)dy.
\end{aligned} \tag{22}
$$

Line (a) is the "summing up" operation of $f_{X,Y}(x,y)$ values for which the condition $z = x + y^2$ is satisfied. In line (b), the order of integration is exchanged, that is, the inner integration is with respect to x after the exchange. Line (c) is due to the evenness of $\delta(x)$. Line (d) is due to the sifting property. Line (e) is the factorization of joint density in terms of the conditional density of X given Y.

Comment

By changing the integration order in line (c), the variable for the function $\delta(z-x-y^2)$ becomes x. After the order change, the variables z and y are treated as constants, and we have the result in line (d). If the inner integral in line (c) were with respect to the variable y, that is, if we do not change the order of integration, we need to use result given in Example 2 to evaluate the integral involving $\delta(z-x-y^2)$.

Example 4

Show that the Fourier transform of $f(t) = 1$ is $F(\Omega) = 2\pi\delta(\Omega)$, where $F(\Omega) = \mathcal{F}\{f(t)\} = \int_{-\infty}^{\infty} f(t)\exp(-j\Omega t)dt$ is the Fourier transformation operation.

Freshman calculus results state that $\int_{-\infty}^{\infty} f(t)\exp(-j\Omega t)dt$ does not converge for any Ω for $f(t) = 1$. Hence, the well-known Fourier transform pair of $1 \leftrightarrow 2\pi\delta(\Omega)$ should be interpreted in the generalized sense. To show $\mathcal{F}\{1\} \overset{(g)}{=} 2\pi\delta(\Omega)$, we need to examine the response of the function $F(\Omega) = \mathcal{F}\{1\}$ to a test function $\Phi(\Omega)$:

$$
\begin{aligned}
\langle \mathcal{F}\{1\}, \Phi(\Omega) \rangle &\overset{(a)}{=} \int_\Omega \mathcal{F}\{1\}\Phi(\Omega)d\Omega \\
&\overset{(b)}{=} \int_\Omega \Big(\int_t 1 e^{-j\Omega t} dt \Big) \Phi(\Omega)d\Omega \\
&\overset{(c)}{=} \int_t \Big(\int_\Omega \Phi(\Omega) e^{-j\Omega t} d\Omega \Big) dt \\
&\overset{(d)}{=} \int_t 2\pi\phi(-t)dt \\
&\overset{(e)}{=} 2\pi\Phi(0) \\
&\overset{(f)}{=} 2\pi \int \delta(\Omega)\Phi(\Omega)d\Omega \\
&\overset{(g)}{=} \langle 2\pi\delta(\Omega), \Phi(\Omega) \rangle .
\end{aligned} \tag{23}
$$

Line (a) is due to the linear functional definition. Line (b) results from the definition of the Fourier transform. Line (c) changes the integration order. Line (d) is due to the inverse Fourier transform relation for ordinary, absolutely integrable functions [1]. [With the assumed test function class (Gaussian functions), the Fourier integral, that is, $F\{\phi(t)\} = \Phi(\Omega)$, is guaranteed to converge in the ordinary calculus sense.] Line (e) is due to the fact that $\Phi(0) = \int_t \phi(t)dt$, that is, the area of the time-domain function, is the value of its Fourier representation at $\Omega = 0$. Lines (f) and (g) are different ways of writing line (e). Considering the leftmost and rightmost sides of (23) and remembering the generalized equality definition in (2), we can conclude the proof of $\mathcal{F}\{1\} \overset{(g)}{=} 2\pi\delta(\Omega)$.

Comment

In a first course, this relation is given by finding the inverse Fourier transform of $2\pi\delta(\Omega)$, i.e., $\mathcal{F}^{-1}\{2\pi\delta(\Omega)\}$, without mentioning the existence of the Fourier integral for $f(t) = 1$. The Fourier integral for $f(t) = 1$ diverges in the usual sense but exists only in the generalized sense or in the sense of distributions.

Example 5

Show that the Fourier transform of $f(t) = \mathrm{sgn}(t)$ is $F(\Omega) \overset{(g)}{=} 2/j\Omega$.

The Fourier transform of $\mathrm{sgn}(t)$,

$$
\mathrm{sgn}(t) = \begin{cases} 1 & t > 0 \\ -1 & t < 0 \end{cases},
$$

can be written as the integral

$$
\mathcal{F}\{\mathrm{sgn}(t)\} = \frac{2}{j} \int_0^{\infty} \sin(\Omega t)dt, \tag{24}
$$

which does not converge in the ordinary calculus sense. Hence, as suspected, $\mathcal{F}\{\mathrm{sgn}(t)\}$ is equal to $2/j\Omega$ in the distribution sense. It is interesting to note that there is no Dirac delta function in the expression $\mathcal{F}\{\mathrm{sgn}(t)\} \overset{(g)}{=} 2/j\Omega$, immediately giving away that the equality is in the generalized sense.

Let's define a regular function $g_T(\Omega)$ as $g_T(\Omega) = \int_0^T \sin(\Omega t)dt = (1-\cos(\Omega T))/\Omega$. We would like to take the limit of $g_T(\Omega)$ as $T \to \infty$ with the goal of evaluating the transform in (24). To do that, we need to examine the response of $g_T(\Omega)$ to a test function $\Phi(\Omega)$, that is, $\langle g_T(\Omega), \Phi(\Omega) \rangle$, and then evaluate the limit of the response as $T \to \infty$.

For a fixed T, $\langle g_T(\Omega), \Phi(\Omega) \rangle$ can be expressed as

$$\langle g_T(\Omega), \Phi(\Omega) \rangle = \langle \frac{1}{\Omega}, \Phi(\Omega) \rangle - \langle \frac{\cos(\Omega T)}{\Omega}, \Phi(\Omega) \rangle = \langle \frac{1}{\Omega}, \Phi(\Omega) \rangle - \langle \cos(\Omega T), \frac{\Phi(\Omega)}{\Omega} \rangle. \quad (25)$$

As $T \to \infty$, the equality in (25) approaches

$$\lim_{T\to\infty} \langle g_T(\Omega), \Phi(\Omega) \rangle = \langle \frac{1}{\Omega}, \Phi(\Omega) \rangle - \lim_{T\to\infty} \langle \cos(\Omega T), \frac{\Phi(\Omega)}{\Omega} \rangle. \quad (26)$$

From (26), it is clear that we need to show $\lim_{T\to\infty} \langle \cos(\Omega T), (\Phi(\Omega)/\Omega) \rangle = 0$ to conclude the proof. Since the test function class is the class of Gaussian functions, the function $\Phi(\Omega)/\Omega$ is absolutely integrable in $\Omega \in (-\infty, \infty)$ in the Cauchy principle value sense. (The Cauchy principle value integral is required due to the singularity of $\Phi(\Omega)/\Omega$ at $\Omega = 0$ [4, p. 359]).

We know from Dirichlet conditions that the Fourier transform of an absolutely integrable function exists in the regular sense [1, p. 290]. An important but less-known fact by the signal processing audience is the Riemann–Lebesgue lemma, stating that, if $x(t)$ is absolutely summable, then $X(\Omega) \to 0$ as $\Omega \to \infty$ [3, p. 278].

Armed with this knowledge, $\langle \cos(\Omega T), (\Phi(\Omega)/\Omega) \rangle$ can be interpreted as the real part of the $\mathcal{F}\{\Phi(\Omega)/\Omega\}$ with the transform-domain variable T. Then, due to the absolute integrability of $\Phi(\Omega)/\Omega$ and the Riemann–Lebesgue lemma, we have $\lim_{T\to\infty} \langle \cos(\Omega T), (\Phi(\Omega)/\Omega) \rangle = 0$.

By multiplying both sides of (26) by $2/j$ and replacing $g_T(\Omega)$ with $\int_0^T \sin(\Omega t)\, dt$, we reach

$$\lim_{T\to\infty} \int_{-\infty}^{\infty} \left(\frac{2}{j} \int_0^T \sin(\Omega t)\, dt \right) \Phi(\Omega)\, d\Omega = \int_{-\infty}^{\infty} \left(\frac{2}{j\Omega} \right) \Phi(\Omega)\, d\Omega, \quad (27)$$

stating that $\mathcal{F}\{\mathrm{sgn}(t)\} \overset{(g)}{=} 2/j\Omega$ via the generalized limit definition given in (10).

Comment

A first course in signal processing needs to sugarcoat some definitions and even some calculations due to pedagogical reasons. Among these, the Fourier transformations of the sign function and unit step function stand out. The sign function, $\mathrm{sgn}(t)$, is clearly not absolutely or square summable; hence, its Fourier transform cannot be given in the usual sense.

In spite of that, to show this result, some instructors calculate the Fourier transform of a regular, absolutely summable function $\mathrm{sgn}(t) e^{-\alpha|t|}$; evaluate the limit of the result as $\alpha \to 0$; and then present the limit as the Fourier transform of $\mathrm{sgn}(t)$. The end result of this calculation matches the correct result, but the intermediate steps, especially the one involving the movement of the limit operation inside of the Fourier transform integral in the final step, are highly questionable. It should be clear at this point that any treatment of integrals diverging in the ordinary calculus sense requires some extraordinary effort. The definition of generalized functions is an effort along this line.

As expected, the Fourier transform of $u(t)$ is also only valid in the generalized sense. By expressing $u(t)$ as $u(t) = (\mathrm{sgn}(t) + 1)/2$ and applying the linearity of the Fourier transform, we can show $\mathcal{F}\{u(t)\} \overset{(g)}{=} 1/j\Omega + \pi\delta(\Omega)$.

Example 6

Find the inverse unilateral Laplace transform of $X(s) = s^2/(s+3)$.

This problem is typically solved by partial fraction expansion, that is,

$$X(s) = \frac{s^2}{s+3} = s - 3 + \frac{9}{s+3}, \quad (28)$$

followed by inverse Laplace transformation via transform-pair recognition. The final answer of this example is $x(t) = \delta^{(1)}(t) - 3\delta(t) + 9\exp(-3t)u(t)$. Our goal is to derive the same result via some alternative paths to illustrate the usage of generalized differentiation.

Let's first express $X(s)$ as $X(s) = s^2 X_p(s)$, where $X_p(s) = 1/(s+3)$. The inverse Laplace transform of $X_p(s)$ is $x_p(t) = \exp(-3t)u(t)$. Hence, the inverse Laplace transform $X(s) = s^2 X_p(s)$ becomes $x(t) = (d^2/dt^2)x_p(t)$. We can verify this result by remembering that the unilateral Laplace transform of $(d/dt)x(t)$ is $sX(s) - x(0^-)$. Note that $x_p(t)$ and its derivatives are all zero at $t = 0^-$ due to the existence of the $u(t)$ term in $x_p(t)$. Let's evaluate the first two derivatives of $x_p(t)$ and compare the result with the answer by partial fraction expansion:

$$\begin{aligned}
x_p^{(1)}(t) &= \frac{d}{dt}\{\exp(-3t)u(t)\} \\
&\overset{(a)}{=} \frac{d}{dt}\{\exp(-3t)\}u(t) + \exp(-3t)\frac{d}{dt}\{u(t)\} \\
&= -3\exp(-3t)u(t) + \exp(-3t)\delta(t) \\
&\overset{(b)}{=} -3\exp(-3t)u(t) + \delta(t), \\
x_p^{(2)}(t) &= \frac{d}{dt} x_p^{(1)}(t) \\
&= \frac{d}{dt}\{-3\exp(-3t)u(t) + \delta(t)\} \\
&\overset{(a)}{=} \frac{d}{dt}\{-3\exp(-3t)\}u(t) - 3\exp(-3t)\frac{d}{dt}\{u(t)\} + \frac{d}{dt}\{\delta(t)\} \\
&= 9\exp(-3t)u(t) - 3\exp(-3t)\delta(t) + \delta^{(1)}(t) \\
&\overset{(b)}{=} 9\exp(-3t)u(t) - 3\delta(t) + \delta^{(1)}(t).
\end{aligned} \quad (29)$$

Line (a) of both equations is due to the product rule for differentiation and the generalized equality of $(d/dt)u(t) \overset{(g)}{=} \delta(t)$. Line (b) is due to the multiplication property of the Dirac delta function from Table 1. Note that the equalities given in (29) are not ordinary equalities but valid only in the generalized sense. The absence of the $\overset{(g)}{=}$ symbol can be a source of inconsistencies and confusion, yet we go back to the conventional notation and symbols in this last example.As a final exercise, let's redo the calculation by evaluating the second derivative of $x_p(t) = f(t)g(t)$ with $f(t) = \exp(-3t)$ and $g(t) = u(t)$

(continued on page 203)

David Ernesto Troncoso Romero, Miriam Guadalupe Cruz Jiménez, and Uwe Meyer-Baese

Alternative Data Paths for the Cascaded Integrator–Comb Decimator

Alternative data paths for the cascaded integrator–comb (CIC) decimator are presented that can be derived upon applying one or two modifications. In the first modification, the integrator–comb pair that surrounds the downsampling register is replaced by a simpler Integrate-and-Dump (ID) block that performs the same decimation task. The resulting data path is referred to as CIC-ID. In the second modification, the remaining comb filters are replaced by a time-multiplexed comb, and the resulting data path is referred to as CIC-ID-TMUX. Compared with a classic N-stage CIC, CIC-ID saves one arithmetic unit and one register, whereas CIC-ID-TMUX uses the same number of registers but trades $(N-1)$ arithmetic units for a simpler two-to-one multiplexer. Both CIC-ID and CIC-ID-TMUX perform slightly fewer arithmetic operations per output sample because of the absence of the comb filter absorbed by the ID block. Pruning formulas for both architectures are presented along with implementation details.

Introduction

Because of its regularity and simple reconfiguration capability, the CIC decimator is employed in many sampling rate-conversion processes, e.g., in multistandard receivers or radio-frequency-to-baseband sigma–delta analog-to-digital converters. The CIC decimator is derived after applying multirate identities to a recursive finite-impulse response filter with transfer function

$$H(z)=\left(\frac{1-z^{-R}}{1-z^{-1}}\right)^{N}, \qquad (1)$$

moving the comb part $(1-z^{-R})^{N}$ to the low-rate section [1]. In (1), R is the integer downsampling factor, and N is the number of CIC pairs.

On the axis of angular frequencies normalized to the sampling rate, ω, the magnitude response $|H(e^{j\omega})|$ obtained from (1) upon replacing $z=e^{j\omega}$ is

$$|H(e^{j\omega})|=\left|\left(\frac{\sin[\omega R/2]}{\sin[\omega/2]}\right)^{N}\right|, \qquad (2)$$

which has useful nulls on multiples of $\omega=2\pi/R$, allowing it to be employed in the first stage of a multistage decimation chain. This is so because, in that stage, whose downsampling factor is R, the decimation filter must have a narrow passband from $\omega=0$ to $\omega=\pi/RP$, and it must have stopbands centered on multiples of $\omega=2\pi/R$, each with a bandwidth of $2\pi/RP$, where P is the product of the downsampling factors of the subsequent stages (filters of the other stages attenuate at the frequency regions not covered by these stopbands).

Because of the shape of $|H(e^{j\omega})|$, there is a droop on the passband, but this can be corrected by other filters in the decimation chain or with specially designed passband-droop compensators, such as the ones described in [2]. On the other hand, the worst-case attenuation of $|H(e^{j\omega})|$ occurs at $\omega=2\pi/R-\pi/RP$, i.e., at the left edge of the first stopband. There are methods to improve that worst-case attenuation, such as those discussed in [3]–[6], that modify $H(z)$ and result in useful CIC-based filters. Yet a practical and prevalent way consists in increasing N accordingly, which preserves the simplicity of the CIC system. Figure 1(a) shows an example of $|H(e^{j\omega})|$ for $R=16$ and $P=4$, where we observe that $|H(e^{j\omega})|$ improves its worst-case attenuation by about 34 dB, from −17 dB to −51 dB, upon changing N from 1 to 3, whereas the passband droop worsens by only about 0.5 dB ($|H(e^{j\omega})|$ is scaled by $1/R^{N}$ to normalize 0 dB at $\omega=0$).

The CIC decimator is used in many scenarios because of its compact architecture, which also has the capability for easy reconfiguration of the downsampling factor R. Figure 1(b) shows the register transfer level (RTL) data path of the CIC decimator (assuming two's complement arithmetic), including input and output interface registers. The data path consists of N integrators working at the high-rate section, where registers remain enabled every clock cycle, and N comb filters working at the low-rate section, where registers are enabled every Rth clock cycle. In that figure, B is the input wordlength and $K=\lceil \log_2(R^{N}) \rceil$

Digital Object Identifier 10.1109/MSP.2021.3052752
Date of current version: 28 April 2021

 1053-5888/21©2021IEEE

is the internal word length growth necessary to prevent overflow in the integrators, as detailed in [1]. Since the value K may be undesirably large, some least significant bits can be truncated at the input of every integrator and comb, which reduces the hardware utilization at the expense of introducing error at the filter output. The controlled way to do this truncation, suggested by Hogenauer in [1], is called *pruning*. Locations for pruning are illustrated in Figure 1(b) with small squares, where B_k denotes the number of pruned bits at the kth position.

If some latency is permitted, throughput can be improved by inserting pipeline registers. In this case, we have to multiply $H(z)$ by a delay term z^{-D}, where D is the sum of the number of pipeline registers placed at the high-rate section, plus R times the number of pipeline registers operating at the lower rate. As an example, Figure 1(c) shows a fully pipelined CIC data path. We can see that pipelining for integrators comes for free, and one extra register is needed per comb stage.

The interest in this article is not modifying $H(z)$ to improve the magnitude response as in the aforementioned methods of [3]–[6]. Instead, we are focused on modifying the RTL data path of the CIC decimator such that it can be implemented with fewer hardware resources, preserving a compact form without needing to increase the processing rate at the high-rate section.

The suggested data paths

The suggested data paths are derived from two basic observations. The first observation is that the classic ID unit, shown on the left side of Figure 2(a), operates in the same way as the single integrator–comb pair shown on the right side of Figure 2(a) (see, for example, [7]). It performs the accumulative summation of a stream of samples and releases the final result after R clock cycles while clearing the accumulation register, thus starting a new accumulation without affecting the input stream. The clear is synchronous; thus, the disciplined synchronous design methodology is not affected because that clear does not involve any combinational loop. Therefore, we can replace the innermost integrator–comb pair of the CIC architecture by a simpler ID unit, eliminating the need for the comb.

The second observation applies for the remaining $(N-1)$ cascaded comb units. Since input and output registers of the cascaded interconnection of comb filters are enabled every Rth clock cycle, there is enough time to multiplex a single comb unit if R is sufficiently larger than $(N-1)$, so we can use a single time-multiplexed comb instead of several combs. We do not use that approach for the integrators because they are already operating at the high-rate section, and the time-multiplexing approach would reduce hardware utilization at the cost of further increasing the rate of operation.

The time-multiplexed comb unit shown on the left side of Figure 2(b), derived after applying the method of [8], has the same impulse response as the cascaded comb units shown on the right side of Figure 2(b) if the multiplexer is switched accordingly and the registers inside the loops are enabled $(N-1)$ times per each enable of the input and

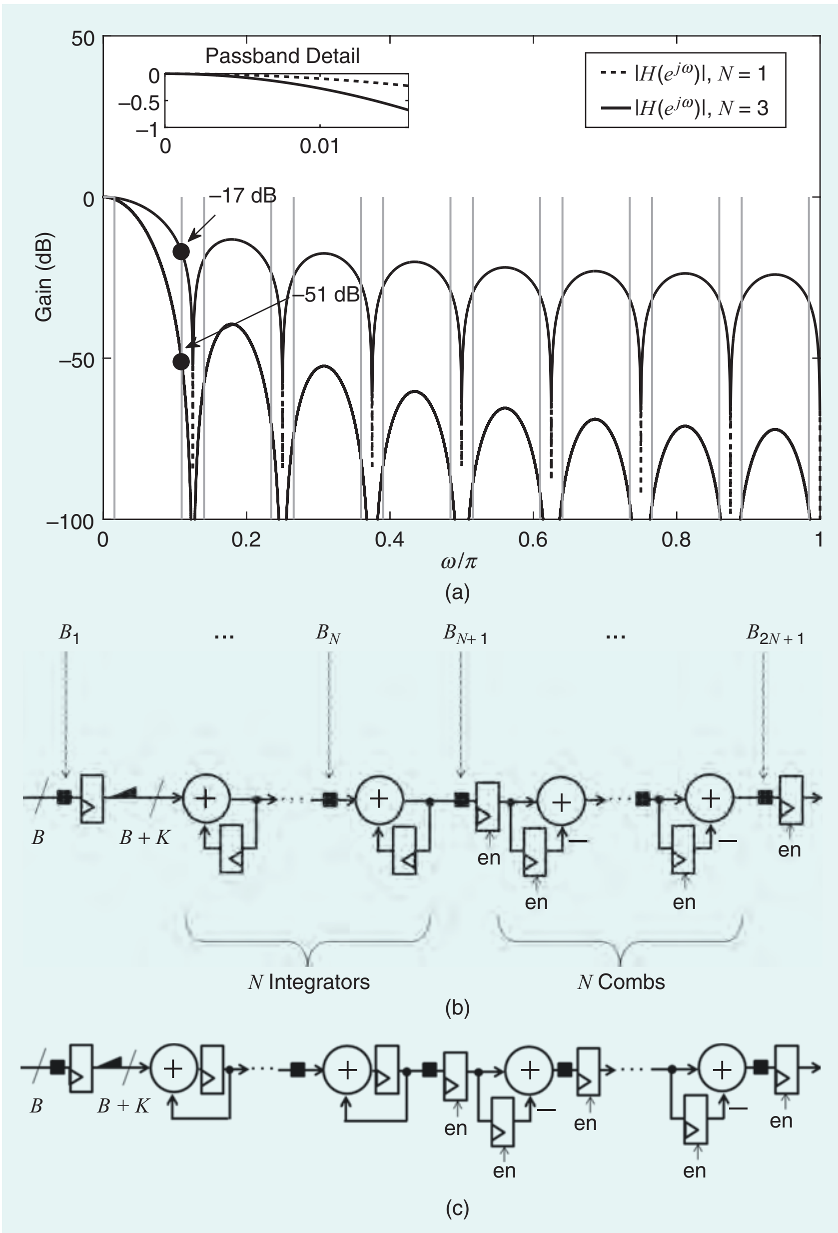

FIGURE 1. CIC filtering: (a) the magnitude response $|H(e^{j\omega})|$ for $N=1$ and $N=3$; (b) the register transfer level (RTL) data path of the CIC decimator; and (c) a fully-pipelined data path. en: enable.

output registers. In other words, the signal enable 2 (“en2”) is asserted every R clock cycles, whereas the signal “en1” is asserted $(N-1)$ times per each assertion of “en2” so that the comb unit performs its $(N-1)$ cascaded operations. The signal “en-mux”’ remains asserted until reaching the first pulse of “en1” to hold the proper routing from the input, and then it is deasserted to route the feedback path.

Figure 2(c) illustrates an example of this timing for a case with $(N-1)=3$ and $R=8$. The control signals “en1” and “en2” are generated at the clock edge = 0 and become effective for their corresponding registers at the clock edge = 1. At that edge, the input register samples new data, and its currently stored data, already processed by the subtractor, are sampled by the subtractor’s register. At the next two pulses of “en1,” the multiplexer routes the feedback path, and the final result of the comb is sampled by the subtractor’s register at the clock edge = 5. After that, the multiplexer routes the data stored in the input register, and the subtractor processes them. At the clock edge = 9, the final result stored in the subtractor’s register is released to the output register, the input register samples new data, and the subtractor’s register samples the value already processed by the subtractor, starting the process again. No matter how we distribute the pulses of “en1” between the pulses of “en2,” the shortest interval between pulses of “en1” lasts m clock cycles, where m is the integer part of $R/(N-1)$. Therefore, to have the time-multiplexed comb unit still operating at a lower rate than that of the integrators, we must have $R \geq m(N-1)$ with at least $m=2$.

The time-multiplexed comb filter has three pipeline registers (one at its input, one at its output, and one at the subtractor’s output). Hence, it introduces an extra delay of R clock cycles in comparison with the cascaded interconnection of comb filters, which has two pipeline registers (one at its input and one at its output).

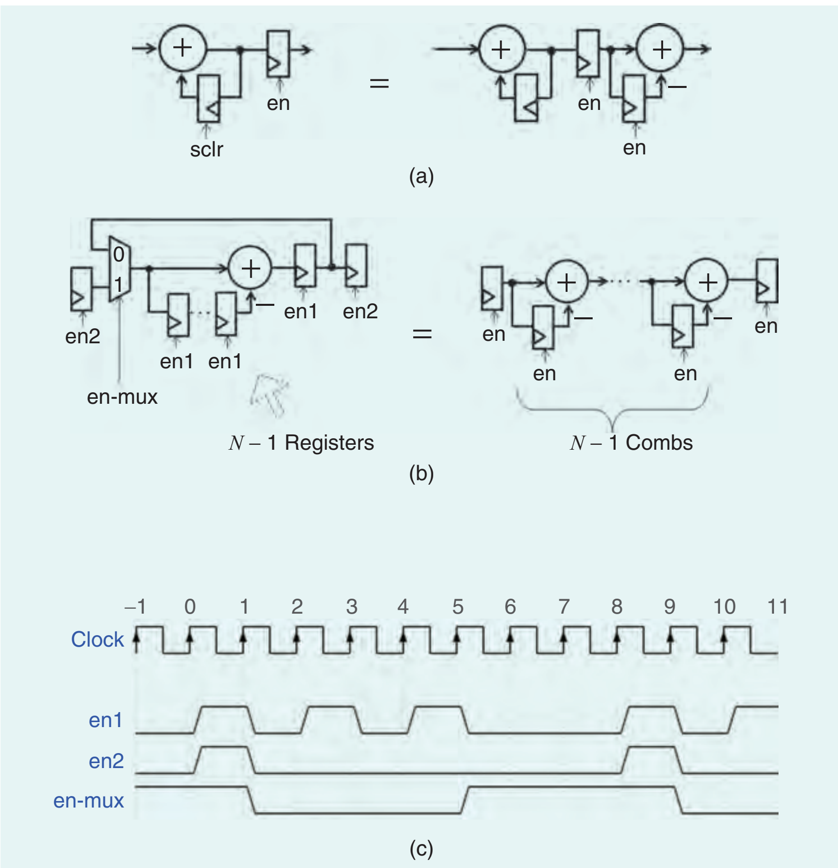

FIGURE 2. Equivalences employed in the suggested data paths: (a) an ID unit and an integrator–comb pair; (b) a time-multiplexed comb unit and a cascade of identical comb units; and (c) a timing example for the time-multiplexed comb unit with $R=8$ and $(N-1)=3$. en: enable; sclr: synchronous clear.

The proposed RTL data paths are obtained by applying either the first one or both of the two preceding observations. The system that results from applying the first observation alone, shown in Figure 3(a) without pipelining, is referred to as CIC-ID. The system that results from applying both observations, shown in Figure 3(b) with pipelined integrators, is referred to as CIC-ID-TMUX. Notice that the storage of the ID block must remain in the feedback path. Again, places for pruning are illustrated with small squares, and B_k denotes the number of pruned bits at the kth position. In both systems, the first $(N-1)$ integrators remain as in the classic architecture, and the innermost integrator–comb pair now becomes the ID block. The other combs, now $(N-1)$ instead of N, are left unrolled in the CIC-ID system, and they are rolled up as a single time-multiplexed unit in CIC-ID-TMUX.

Pruning in the suggested data paths

Equations to compute B_k for the CIC-ID and CIC-ID-TMUX architectures are presented next. These formulas, developed in “Obtaining Equations for Pruning,” follow the same criterion employed by Hogenauer in [1], namely, that the variance of the error due to pruning inside the data path must remain bounded by the variance due to output pruning. In these formulas, F_k^2 denotes the “variance error gain” of the kth error source, computed as the sum of squared impulse response coefficients of the filter seen by that source in front of it (for sources at the high-rate section, the downsampler must be placed at the output of that filter via multirate identities). Note that F_k denotes the square root of F_k^2. These formulas are also given in terms of B_{2N}, the number of bits pruned at the output of the system, which is known beforehand.

In the CIC-ID architecture, we can compute B_k as

$$B_k = \lfloor B_{2N} + \tfrac{1}{2}\log_2 \tfrac{1}{2N-1} - \log_2 F_k \rfloor \quad (3)$$

for $k=1,2,\ldots,2N-1$. In the CIC-ID-TMUX architecture, $B_1, B_2, \ldots, B_N$ can be computed with (3), but for the time-multiplexed comb unit, we have $B_{N+1} = B_{N+2} = \ldots = B_{2N-1}$. Hence, we only need to compute B_{N+1} additionally. We use

$$B_{N+1} = \lfloor B_{2N} + \tfrac{1}{2}\log_2 \tfrac{N-1}{2N-1} - \log_2 F_c \rfloor, \quad (4)$$

where F_c is the square root of F_c^2, which in turn is defined as the collective variance error gain of the comb section, given by

$$F_c^2 = \sum_{k=N+1}^{2N-1} F_k^2. \quad (5)$$

Table 1 intuitively contrasts the numbers of pruned bits B_k in the proposed data paths and in the classic CIC data path for three values of R considering a 16-bit input and 16-bit output, with $N = 5$ cascaded stages. They are grouped as pruning for integrators (PI), pruning for combs (PC), and pruning for output (PO). We observe that B_k remains practically the same for the integrators in all of the architectures. In the CIC-ID architecture, the number of pruned bits for the first comb block is the number of pruned bits for the second comb block of the CIC. This occurs because the first comb unit of the CIC is no longer needed, as it is absorbed by the ID unit. In CIC-ID-TMUX, the number of pruned bits of the time-multiplexed comb is the number of pruned bits of the comb block placed at the middle of the cascade in the CIC. This occurs because the pruning of the time-multiplexed comb is the same in the whole cascade of combs, and its contribution to the output error is equally divided among all of the comb stages.

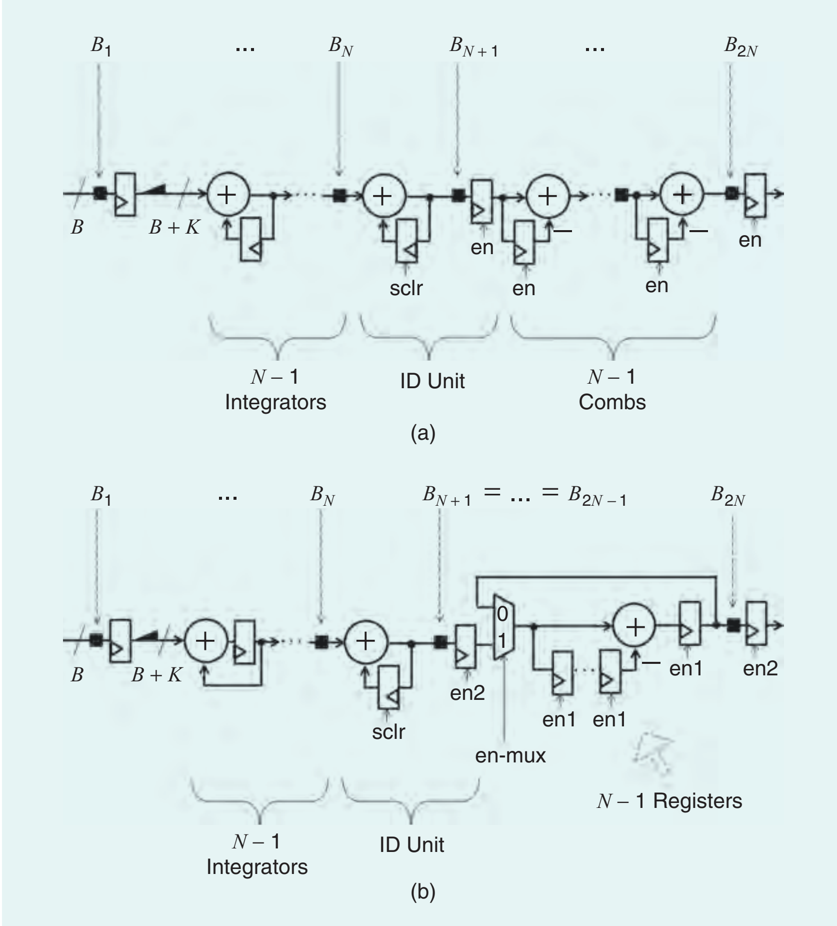

FIGURE 3. The proposed RTL data paths of CIC decimation: (a) CIC-ID. (b) CIC-ID-TMUX. en: enable; sclr: synchronous clear.

Advantages and limitations

The main advantage of the proposed data paths is the reduction of utilized hardware. Table 2 summarizes the hardware utilization and computational complexity for the classic CIC, CIC-ID, and CIC-ID-TMUX data paths. Compared with the classic CIC, CIC-ID-TMUX preserves the same number of registers but trades $(N-1)$ subtractors for a simpler two-to-one multiplexer, whereas CIC-ID saves one subtractor and one register. Both architectures perform a slightly lower amount of arithmetic operations per output sample because of the absence of the comb absorbed by the ID unit.

Table 3 presents the estimated chip area utilization of these data paths, given by (6)–(8) shown at the bottom of this page.

In (6)–(8), W_k is the bus width at the kth pruning position, and A_{mux}, A_{fa}, and A_{ff} are the respective VLSI-Technology Silicon Compiler area costs for a single-bit 2:1 multiplexer, full adder, and flip-flop, which, for a 1-μm CMOS process, are $A_{mux} = 0.0012$ mm^2, $A_{fa} = 0.0086$ mm^2, and $A_{ff} = 0.0037$ mm^2 according to Meyer-Baese [9]. We have considered $N = 5$ stages and a large value $R = 1{,}024$, as in the CIC system included in the decimator HSP43220 by

$$A_{\mathrm{CIC}} = \underbrace{A_{ff} \times W_1}_{\text{input reg.}} + \underbrace{A_{ff} \times W_{N+1}}_{\text{downsamp. reg.}} + \underbrace{A_{ff} \times W_{2N+1}}_{\text{output reg.}} + \underbrace{\sum_{k=1}^{2N} (A_{fa} + A_{ff}) \times W_k}_{\text{integrators and combs}}, \quad (6)$$

$$A_{\mathrm{CIC-ID}} = \underbrace{A_{ff} \times W_1}_{\text{input reg.}} + \underbrace{A_{ff} \times W_{N+1}}_{\text{downsamp. reg.}} + \underbrace{A_{ff} \times W_{2N}}_{\text{output reg.}} + \underbrace{\sum_{k=1}^{2N-1} (A_{fa} + A_{ff}) \times W_k}_{\text{integrators and combs}}, \quad (7)$$

$$A_{\mathrm{CIC-ID-TMUX}} = \underbrace{A_{ff} \times W_1}_{\text{input reg.}} + \underbrace{A_{ff} \times W_{N+1}}_{\text{downsamp. reg.}} + \underbrace{\sum_{k=1}^{N} (A_{fa} + A_{ff}) \times W_k}_{\text{integrators}}$$
$$+ \underbrace{(A_{mux} + A_{fa} + N \times A_{ff}) \times W_{N+1}}_{\text{comb}} + \underbrace{A_{ff} \times W_{2N}}_{\text{output reg.}}. \quad (8)$$

Obtaining Equations for Pruning

Consider that the error at the kth noise source due to pruning has a uniform probability distribution with a width of $E_k = 2^{B_k}$ Thus, the variance of the error is $\sigma_k^2 = E_k^2/12 = 2^{2B_k}/12$. The total variance contributed by the kth error source is $\sigma_{T,k}^2 = \sigma_k^2 F_k^2$, i.e.,

$$\sigma_{T,k}^2 = (\tfrac{1}{12} 2^{2B_k}) F_k^2, \tag{S1}$$

where F_k^2 is the variance error gain for the kth error source. The total variance, σ_T^2, is

$$\sigma_T^2 = \sum_{k=1}^{2N} \sigma_{T,k}^2 = \underbrace{\sum_{k=1}^{N} \sigma_{T,k}^2}_{\text{integrators' sources}} + \underbrace{\sum_{k=N+1}^{2N-1} \sigma_{T,k}^2}_{\text{combs' sources}} + \underbrace{\sigma_{T,2N}^2}_{\text{output source}}. \tag{S2}$$

Letting the total variances of every error source from $k=1$ to $k=2N-1$ contribute at most $1/(2N-1)$ of the variance at the output source, we have

$$\sigma_{T,k}^2 \le \tfrac{1}{2N-1} \sigma_{T,2N}^2 \quad \text{for} \quad k = 1, 2, \ldots, 2N-1. \tag{S3}$$

Using (S1) in (S3) with the consideration that the output source error sees a unitary gain in front of it, i.e., $F_{2N}^2 = 1$, we obtain

$$(\tfrac{1}{12} 2^{2B_k}) F_k^2 \le \tfrac{1}{2N-1} (\tfrac{1}{12} 2^{2B_{2N}}) \quad \text{for} \quad k = 1, 2, \ldots, 2N-1 \tag{S4}$$

We can easily derive (3) from (S4).

For the time-multiplexed comb unit, we have $B_{N+1} = B_{N+2} = \cdots = B_{2N-1}$. Hence, we can compute B_k for $k = 1, 2, \ldots, N$ using (3), and we only need to compute B_{N+1} additionally. For the $(N-1)$ error sources at the comb section, grouped in (S2) as "combs' sources," we define

$$\sigma_c^2 = \sum_{k=N+1}^{2N-1} \sigma_{T,k}^2. \tag{S5}$$

Applying the constraint (S3) to the collective variance of the comb section, we can write

$$\sigma_c^2 \le \tfrac{N-1}{2N-1} \sigma_{T,2N}^2. \tag{S6}$$

Using (S1) in (S6) and considering $B_{N+1} = B_{N+2} = \cdots = B_{2N-1}$, we can write

$$\sigma_c^2 = (\tfrac{1}{12} 2^{2B_{N+1}}) F_c^2, \tag{S7}$$

where F_c^2 is given in (5). Using (S7) in (S6) and also using (S1) in (S6) with the consideration that $F_{2N}^2 = 1$, we have

$$(\tfrac{1}{12} 2^{2B_{N+1}}) F_c^2 \le \tfrac{N-1}{2N-1} (\tfrac{1}{12} 2^{2B_{2N}}). \tag{S8}$$

We can easily derive (4) from (S8).

Table 1. A comparison of the number of pruned bits for $N = 5$, a 16-bit input, and a 16-bit output.

	CIC	CIC-ID	CIC-ID-TMUX
R = 8	PI: {0, 3, 5, 7, 8} PC: {9, 10, 11, 12, 12} PO: {15}	PI: {0, 3, 5, 7, 8} PC: {10, 11, 12, 12} PO: {15}	PI: {0, 3, 5, 7, 8} PC: {11} PO: {15}
R = 128	PI: {2, 9, 15, 21, 26} PC: {29, 30, 31, 32, 32} PO: {35}	PI: {2, 9, 15, 21, 26} PC: {30, 31, 32, 32} PO: {35}	PI: {2, 9, 15, 21, 26} PC: {31} PO: {35}
R = 1,024	PI: {3, 13, 23, 31, 40} PC: {44, 45, 46, 47, 47} PO: {50}	PI: {4, 13, 23, 31, 40} PC: {45, 46, 47, 47} PO: {50}	PI: {4, 13, 23, 31, 40} PC: {46} PO: {50}

PI: pruning for integrators; PC: pruning for combs; PO: pruning for output.

Intersil [10]. Additionally, we included an example with $N = 5$ stages and $R = 8$, which is the smallest possible value of R for which the time-multiplexed comb still operates at a lower rate than the integrators. For both examples, we consider a 16-bit input and 16-bit output. In the CIC-ID-TMUX system, there must be an additional control for the time-multiplexed comb, but it only needs a little extra logic added to the overall decimation control and does not involve important complexity in a decimation chain where the data path consumes most of the hardware utilization. Thus, complexity of control logic is not included. We observe savings between 7 and 10% for the CIC-ID system in comparison to CIC, and savings between 20 and 30% for CIC-ID-TMUX in comparison to CIC.

We have also mapped the CIC, CIC-ID, and CIC-ID-TMUX architectures to the field-programmable gate array (FPGA) Intel chip EP4CE115F29C7, popular among the academic community (available in the kit DE2-115), using the aforementioned values for R, N, and input/output wordlengths. In all cases, pipelined integrators were employed (except for the ID block). We used Quartus Prime 18.0 Lite as a synthesis tool, and for static timing analysis we employed the Quartus Prime Timing Analyzer under the 85 C model (the worst-case scenario). For power consumption estimation, we simulated a test signal passing through these systems and used the Quartus Prime Power Analyzer to perform that estimation with the switching activity data generated from the simulations. The test signal, a chirp with linearly increasing frequency from 0 to the Nyquist limit, was provided to the systems for 199 µs. Simulations were done with ModelSim Intel FPGA Starter Edition 10.5 b, and the systems were clocked at 50 MHz, a clock frequency supported by all of them.

Table 4 summarizes the hardware cost in logic elements (LEs), estimation of maximum frequency of operation, estimation of power consumption, and latency in the system's response. We observe a reduction in the number of LEs in the proposed architectures. In particular, the CIC-ID-TMUX system consumes some bits of embedded random-access memory (RAM) because the chain of registers of its time-multiplexed comb is mapped to this resource. As a consequence, fewer LEs are required in comparison to the other systems. Besides, we observe a greater latency for that system because of the inherent pipelining of its comb section, and higher power consumption because of the higher switching rate of the time-multiplexed operation. The CIC-ID system also has a reduced number of LEs and does not use RAM resources. However, it cannot take full advantage of pipelined integrators because the ID block must remain nonpipelined. Therefore, its maximum frequency of operation is slightly lower, and its power consumption may grow if the data path bus becomes too large (which makes the ID block too wide).

In general, the proposed approach has the following limitations:

- The storage of the ID block should remain in the recursive path to preserve the equivalence illustrated in Figure 2(a). This can make that block slower in comparison with a pipelined integrator. If that periodically cleared accumulation is implemented in pipelined form, the data path would correspond to a decimator whose filter's transfer function is $H(z) = [(1-z^{-R})/(1-z^{-1})]^{N-1} \times [(1-z^{-(R-1)})/(1-z^{-1})]$, which slightly decreases the worst-case attenuation value.
- The CIC-ID-TMUX data path is not efficient if $R < 2(N-1)$ holds because, in that case, the processing rate of the time-multiplexed comb unit would need to increase and become equal to or higher than the processing rate of the high-rate section. In general, the combinational delay of the time-multiplexed comb unit must be lower than the shortest interval between the pulses of "en1," a constraint that can be met more easily as R grows; i.e., as m grows in the relation $R \geq m(N-1)$. The switching activity of this block might also make it more power consuming in comparison with the cascaded combs.
- The time-multiplexed comb filter introduces an extra latency of R clock cycles for the system's response in comparison with the cascaded interconnection of nonpipelined comb filters. The CIC-ID-TMUX data path is not efficient if latency is not tolerated.
- Implementation of the CIC-ID-TMUX data path on FPGAs may not see a clear benefit. The reason is that, in FPGAs, a basic building block includes both combinational logic and registers. Hence, the implementation of a single register needs approximately the same number of basic building blocks than the implementation of an arithmetic unit plus a register. In this sense, there is no advantage in the fact that

Table 2. A comparison of computational complexity and hardware utilization of classic and proposed data paths for the CIC decimator.

	Number of Arithmetic Operations per Output Sample	Number Of Arithmetic Units	Number of Registers	Two-to-One Multiplexer
Classic CIC	$N \times R + N$	$2N$	$2N+3$	0
CIC-ID	$N \times R + N - 1$	$2N-1$	$2N+2$	0
CIC-ID-TMUX	$N \times R + N - 1$	$N+1$	$2N+3$	1

Table 3. A comparison of the estimated chip area using silicon compiler data for a 1-μm CMOS process.

	CIC	CIC-ID	CIC-ID-TMUX
		No Pruning	
$R = 8$	4.05 mm^2	3.67 mm^2	3.02 mm^2
$R = 1{,}024$	8.49 mm^2	7.67 mm^2	6.29 mm^2
		Pruning	
$R = 8$	3.07 mm^2	2.8 mm^2	2.39 mm^2
$R = 1{,}024$	4.14 mm^2	3.85 mm^2	3.44 mm^2

Table 4. Comparison of synthesis results on the EP4CE115F29C7 FPGA chip.

	CIC	CIC-ID	CIC-ID-TMUX
		No Pruning	
$R = 8$	LEs: 395 MFO: 98.33 MHz PwC: 147.03 mW L: 0.350 µs	LEs: 364 MFO: 97.88 MHz PwC: 149.3 mW L: 0.350 µs	LEs: 275 + 62 memory bits MFO: 112.74 MHz PwC: 154.59 mW L: 0.490 µs
$R = 1{,}024$	LEs: 828 MFO: 87.7 MHz PwC: 146.89 mW L: 40.99 µs	LEs: 762 MFO: 76.3 MHz PwC: 157.32 mW L: 40.99 µs	LEs: 591 + 132 memory bits MFO: 83.51 MHz PwC: 158.93 mW L: 61.450 µs
		Pruning	
$R = 8$	LEs: 298 MFO: 98.47 MHz PwC: 153.72 mW L: 0.350 µs	LEs: 276 MFO: 98.33 MHz PwC: 143.76 mW L: 0.350 µs	LEs: 223 + 40 memory bits MFO: 112.71 MHz PwC: 153.30 mW L: 0.490 µs
$R = 1{,}024$	LEs: 398 MFO: 90.88 MHz PwC: 153.52 mW L: 40.99 µs	LEs: 378 MFO: 83.89 MHz PwC: 141.58 mW L: 40.99 µs	LEs: 341 + 40 memory bits MFO: 74.45 MHz PwC: 154.71 mW L: 61.450 µs

LE: logic element; MFO: maximum frequency of operation; PwC: power consumption: L: latency.

the CIC-ID-TMUX architecture uses fewer arithmetic units than the classic architecture because the number of registers is the same in both data paths.

Closing comments

We have observed that in the classic CIC decimator the comb units do not need to run concurrently. On the one hand, the innermost integrator–comb pair is the same as an ID unit; thus, we can eliminate the comb in that pair and replace its function by a much simpler synchronous clear, leading to the CIC-ID architecture. The other combs are necessary, but since they operate at the low-rate section, time multiplexing does not come at the cost of having to increase the maximum rate of operation. Hence, we can replace them by a simpler time-multiplexed comb unit, leading to the CIC-ID-TMUX architecture. The proposed data paths for the CIC decimator can achieve, in comparison with the traditional CIC data path, lower hardware utilization without compromising computational complexity. This makes them useful options over the classic CIC architecture.

Authors

David Ernesto Troncoso Romero (detroncosoro@conacyt.mx) is a researcher at Consejo Nacional de Ciencia y Tecnología (CONACYT), México, commissioned to Escuela Superior de Cómputo, Instituto Politécnico Nacional, México City, 07738, México, as a "Catedrático CONACYT" visiting professor. His research interests include the design and implementation of low-complexity digital signal processing algorithms.

Miriam Guadalupe Cruz Jiménez (miriam.gcj@gmail.com) is a postdoctoral researcher at Centro de Investigación en Computación, Instituto Politécnico Nacional, México City, 07738, México. Her research interests include digital filters, filter banks, and multirate systems as well as their efficient implementations.

Uwe Meyer-Baese (uwe.meyer-baese@ieee.org) received his MS, Dr.-Ing. "Summa cum Laude," and Dr.-Ing. (Habil) from the Darmstadt University of Technology, Darmstadt, Germany, in 1989, 1995, and 2003, respectively. He is currently a full professor in the Department of Electrical and Computer Engineering at Florida State University, Tallahassee, Florida, 32310, USA. He holds three patents, has published more than 130 journal and conference papers and seven books, and has supervised more than 60 master and Ph.D. thesis projects. He is the author of best-selling textbooks on digital signal processing with field-programmable gate arrays (Springer-Verlag). He received the Max-Kade Award in Neuroengineering in 1997 and the Humboldt Research Award in 2009.

References

[1] E. B. Hogenauer, "An economical class of digital filters for decimation and interpolation," *IEEE Trans. Acoust., Speech and Signal Process.*, vol. 29, no. 2, pp. 155–162, 1981. doi: 10.1109/TASSP.1981.1163535.

[2] D. E. T. Romero and M. G. C. Jimenez, "Efficient wide-band droop compensation for CIC filters: Ad-hoc and reconfigurable FIR architectures," *Electron. Lett.*, vol. 53, no. 4, pp. 228–229, 2017. doi: 10.1049/el.2016.3782.

[3] T. Saramaki and T. Ritoniemi, "A modified comb filter structure for decimation," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, 1997, pp. 2353–2356.

[4] J. O. Coleman, "Chebyshev stopbands for CIC decimation and CIC-implemented array tapers in 1D and 2D," *IEEE Trans. Circuits Syst. I – Reg. Papers*, vol. 59, no. 12, pp. 2956–2968, 2012. doi: 10.1109/TCSI.2012.2206435.

[5] D. E. T. Romero, "Efficient CIC-based architecture with improved aliasing rejection and reduced computational complexity," *Electron. Lett.*, vol. 52, no. 15, pp. 1294–1295, 2016. doi: 10.1049/el.2016.1672.

[6] M. G. C. Jimenez, U. Meyer-Baese, and G. Dolecek, "Computationally-efficient CIC-based filter with embedded Chebyshev sharpening for the improvement of aliasing rejection," *Electron. Lett.*, vol. 53, no. 4, pp. 281–283, 2017. doi: 10.1049/el.2016.3309.

[7] "Integrate and Dump MATLAB's block for communications," Introduced before R2006a, MATLAB. Accessed: Dec. 13, 2020. [Online]. Available: https://www.mathworks.com/help/comm/ref/integrateanddump.html

[8] Z. Jiang and A. N. Willson Jr., "Efficient digital filtering architectures using pipelining/interleaving," *IEEE Trans. on Circuits and Syst. II – Analog and Digital Signal Processing*, vol. 44, no. 2, pp. 110–119, 1997. doi: 10.1109/82.554438.

[9] U. Meyer-Baese, *Fast Digital Signal Processing (in German).* Heidelberg: Springer-Verlag, 1999.

[10] "HSP43220 Decimating Digital Filter," Intersil, Milpitas, CA, Data sheet, 2008. Accessed: Dec. 13, 2020. [Online]. Available: https://media.digikey.com/pdf/Data%20Sheets/Intersil%20PDFs/hsp43220.pdf

SP

Athina Petropulu

IEEE Signal Processing Society PROGRESS

Support for underrepresented talent in the field of signal processing

Promoting Diversity in Signal Processing (PROGRESS) is a new initiative of the IEEE Signal Processing Society (SPS), aiming to motivate and support women and underrepresented minorities to pursue academic careers in signal processing; see ieeeprogress.org. The PROGRESS logo, as seen in Figure 1, was created by Marija Iloska, a Ph.D. student, and Prof. Petar Djuric, of Stony Brook University, New York, United States. PROGRESS includes workshops at ICASSP and ICIP conferences as well as follow-up surveys and mentoring teleconferences. This article describes the first PROGRESS workshop and reports on the results of a survey that was taken by the workshop participants.

Women and underrepresented minorities account for only 11% of the SPS's membership. At PROGRESS, we believe that diversity and inclusion are pillars of innovation. The key to increasing diversity and inclusion in signal processing is a more diverse faculty. Such faculty offer role models and are well positioned to draw women and underrepresented minorities to the field and inspire them for excellence.

The first PROGRESS workshop was held virtually on 26–27 October 2020, during ICIP 2020. It was attended by 202 students from 34 different countries:

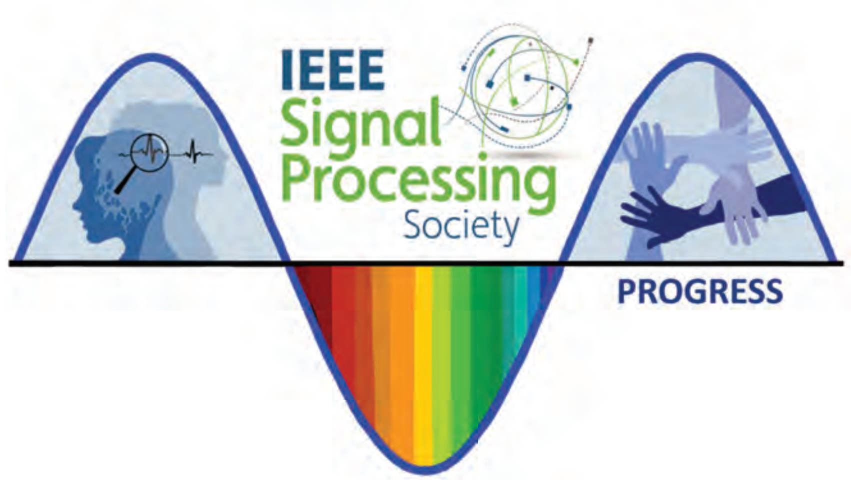

FIGURE 1. The PROGRESS logo; see ieeeprogress.org for more information.

the United States (57), India (42), China (40), Canada (7), the United Arab Emirates (UAE, 6), Bangladesh (5), Spain (3), Sri Lanka (3), Switzerland (2), Tunisia (2), Malaysia (2), the United Kingdom (2), Egypt (2), Ireland (2), Brazil (2), Austria (2), and one each from Colombia, Ecuador, Algeria, Australia, Italy, Lebanon, Indonesia, Ghana, Greece, Pakistan, Peru, Poland, Portugal, Qatar, Saudi Arabia, and Turkey.

The workshop started with a greeting from SPS President Dr. Ahmed Tewfik. Day 1 included an inspirational talk by Prof. Mary (Missy) Cummings, faculty of electrical and computer engineering (ECE) at Duke University, who is a leading researcher in human–robot interaction and one of the U.S. Navy's first female fighter pilots. It also included two panels of high-profile academics from around the world discussing the faculty-hiring process in their countries. The first panel was scheduled at a time convenient for the Eastern Hemisphere. It was moderated by IEEE SPS Vice President, Membership Prof. K.V.S. Hari of the Indian Institute of Science Bangalore, India, and it featured Dr. Lina Karam, dean of the School of Engineering at Lebanese American University, Beirut, Lebanon; Dr. Hussain Al-Ahmad, dean of Engineering at the University of Dubai, UAE; Dr. Zhi-Quan Tom Luo, academic vice president of the Chinese

Digital Object Identifier 10.1109/MSP.2021.3067588
Date of current version: 28 April 2021

University of Hong Kong, Shenzhen, China; and Dr. Markus Rupp, university professor, Technische Universität Wien (TU Wien), Vienna, Austria.

The second panel was scheduled at a time convenient for the Western Hemisphere. It was moderated by Prof. Ana Isabel Pérez-Neira of Universitat Politècnica de Catalunya, Barcelona, Spain, a member of the SPS Board of Governors. It featured Prof. Stella Batalama, dean of the College of Engineering and Computer Science, Florida Atlantic University, Boca Raton, Florida, United States; Prof. Christian Jutten, University Grenoble-Alpes, Grenoble, France; Prof. Konstantina Nikita, National Technical University of Athens, Greece; Prof. Nikolaos Sidiropoulos, chair of ECE at the University of Virginia, Charlottesville, Virgina, United States; Prof. Roy Yates, chair of the Faculty Search Committee of Rutgers University, Piscataway, New Jersey, United States; and Prof. Anderson Rocha, director of the Institute of Computing, University of Campinas, Brazil.

PROGRESS is a new initiative of the IEEE Signal Processing Society, aiming to motivate and support women and underrepresented minorities to pursue academic careers in signal processing.

Day 1 also included two panels of academics sharing their faculty experiences. The first panel, targeting the Eastern Hemisphere, featured Prof. Deepa Kundur, chair of the Department of ECE, University of Toronto, Canada; Prof. Anubha Gupta, Indraprastha Institute of Information Technology, Delhi, India; Prof. Qian He, University of Electronic Science and Technology, Chengdu, China; Prof. Islem Rekik, Istanbul Technical University, Turkey; and Prof. Odette Scharenborg, Delft University of Technology, The Netherlands.

The second panel, targeting the Western Hemisphere, featured Prof. Carol Y. Espy-Wilson of the University of Maryland, College Park, United States; Prof. Piya Pal of the University of California San Diego; Prof. Roxana Saint-Nom of Universidad Argentina de la Empresa, Buenos Aires, Argentina; Prof. Xiangnan Zhong of Florida Atlantic University, United States; and Prof. Donald S. Williamson of Indiana University, Bloomington, United States.

There were also two question and answer (Q&A) sessions, moderated by Prof. Petar Djuric, chair of the Department of ECE at Stony Brook University, New York, United States; Prof. Pascale Fung of the Hong Kong University of Science and Technology, Hong Kong; Prof. Monica Bugallo, director of the Women in Science and Engineering Program at Stony Brook University, New York, United States; Prof. Stella Batalama, dean of the College of Engineering and Computer Science, Florida Atlantic University, United States; Prof. Rabab K. Ward of the University of British Columbia, Vancouver, Canada, who is director of IEEE Division 1X and past president of the SPS; Prof. Raquel Assis of Florida Atlantic University, United States; Prof. Behnaz Ghoraani of Florida Atlantic University, United States; Prof. Xiangnan Zhong of Florida Atlantic University; and Prof. Sareh Taebi of Florida Atlantic University, United States.

Day 2 included professional training on concepts and tools to support career success, including networking and general information from the negotiation literature, delivered by C.K. Gunsalus & Associates. There were also two panels on how one can obtain funding. The participants of the first panel included Prof. Raed Shubair of New York University Abu Dhabi, UAE; SPS Regional Director Prof. Woon-Seng Gan of Nanyang Technological University, Singapore; Prof. Christian Jutten of the University Grenoble-Alpes, France; Prof. Deepa Kundur, chair of the Department of ECE at the University of Toronto, Canada; and Prof. Zhi-Quan Tom Luo, academic vice president of the Chinese University of Hong Kong, Shenzhen, China. The second panel included Prof. Zhi Tian of George Mason University, Fairfax, Virginia, United States, who is also a member of the SPS Board of Governors; Prof. Konstantina Nikita, National Technical University of Athens, Greece; Prof. Ana Isabel Pérez-Neira of Universitat Politècnica de Catalunya, Spain; and Prof. Markus Rupp, TU Wien, Austria.

Talks on advanced educational tools were given by Laura Acion, associate researcher scientist, Instituto de Cálculo, Universidad de Buenos Aires–National Research Council of Argentina, Argentina; Prof. Waheed Bajwa of Rutgers University, Piscataway, New Jersey, United States; and Prof. Raj Rao Nadakuditi of the University of Michigan, Ann Arbor, United States.

A survey was distributed to the students three weeks after the workshop and concluded on 15 December 2020. The survey asked the participants to rate their interest in the various sessions of the workshop. It also asked the participants to rate from 1 to 10 their interest in a postdoctoral or faculty position before attending PROGRESS and after attending PROGRESS. Twenty-five students responded to the survey. The results suggested that, before PROGRESS, 10 of the 25 students expressed interest 8 or higher, while after PROGRESS, 18 students indicated interest 8 or higher. This shift of interest in favor of pursuing an academic position was very encouraging. Also supportive were the comments of the participants, some of which are given here.

- "I found learning from global and international leaders in academia very helpful."
- "The workshop made me aware of various opportunities in foreign countries and the benefits of choosing the field of signal processing."
- "It gave me a better understanding of academia. I could see how faculty members and researchers with academic positions around the world contribute to the broader spectrum of learning (and teaching) and how efforts are being made for inclusivity and diversity. I think it gave me a sense of optimism and confidence to pursue an academic position in the future."

- "I am doing my Ph.D. in a relatively new university. The PROGRESS workshop helped me to get a better exposure of how things operate in other universities and their culture. It really got me motivated when professors of high reputation spent their time to interact and share their knowledge with early researchers like me."
- "In the panel with faculty who shared their experiences, it became clear that everybody struggles at some time in their academic career. This made me more confident that an academic position is actually within my possibilities."
- "The PROGRESS motivated me a lot, especially because we had a lot of wonderful examples of how a career could be merged and coexist perfectly with the private life of anybody and it should be up to us to decide where is the boundary."
- "Hearing how women have actually been able to combine a career in academia and still have a family is very helpful. In my country, there are almost no women in my field of research in permanent academic positions, so there are not really any role models, and it was very interesting to hear from women around the globe about their experiences."

Via the survvey, the students also suggested topics for future PROGRESS workshops, including a session on preparation of a CV, cover letters, statements, grant writing, a list of opportunities (postdoctoral, faculty, and scholarships), a list of platforms where one could find tools to sharpen signal processing skills, a mentorship program, and a forum for Q&A beyond the workshop.

The next PROGRESS workshop will be virtual and is scheduled for 4–5 June 2021—right before ICASSP 2021. More information can be found at ieeeprogress.org.

Author

Athina Petropulu (athinap@soe.rutgers.edu) is a Distinguished Professor of electrical and computer engineering at Rutgers, the State University of New Jersey, president-elect of the IEEE Signal Processing Society, and director of PROGRESS.

SP

LECTURE NOTES

(continued from page 193)

via the Leibniz generalized product rule, $(d^n/dt^n)\{f(t)g(t)\} = \Sigma_{k=0}^{n}\binom{n}{k}f^{(n-k)}(t)\, g^{(k)}(t)$:

$$
\begin{aligned}
x_p^{(2)}(t) &= f^{(2)}(t)g(t) + 2f^{(1)}(t)g^{(1)}(t) \\
&\quad + f(t)g^{(2)}(t) \\
&= 9\exp(-3t)u(t) + 2(-3\exp \\
&\quad (-3t))\delta(t) + \exp(-3t)\delta^{(1)}(t) \\
&\overset{(a)}{=} 9\exp(-3t)u(t) - 6\delta(t) \\
&\quad + [\delta^{(1)}(t) + 3\delta(t)] \\
&= 9\exp(-3t)u(t) - 3\delta(t) \\
&\quad + \delta^{(1)}(t). \qquad (30)
\end{aligned}
$$

In line (a), the basic and advanced versions of the product rule in Table 1 are applied. The advanced product rule states that $f(t)\delta^{(1)}(t) = f(0)\delta^{(1)}(t) - f^{(1)}(0)\delta(t)$, and substituting $f(t) = \exp(-3t)$ into this relation gives the term in the square brackets of line (a). We see that the final result given by either (29) or (30) matches the one by the partial fraction expansion, provided that we handle the differentiation of $x_p(t)$ in the generalized sense, obeying the rules of Dirac delta function manipulation.

What we have learned

We have studied generalized functions, limits, and derivatives as well as their applications in some signal processing problems. These notes aim to show that many familiar equalities are valid only in the generalized sense. Hence, the equality signs should be replaced with $\overset{(g)}{=}$ in many calculations involving Dirac delta functions, unit step functions, and so on. Interested readers can examine classical signal processing textbooks of Papoulis [3] and Bracewell [4] for a brief treatment of generalized functions. For more information, readers are invited to examine [7], [9], and [10].

Acknowledgments

I would like to thank Prof. Bülent Sankur of Boğaziçi University, Istanbul, Turkey, for many suggestions and comments.

Author

Çağatay Candan (ccandan@metu.edu.tr) received his Ph.D. degree in electrical engineering from the Georgia Institute of Technology, Atlanta, in 2004. He is a professor in the Department of Electrical and Electronics Engineering, Middle East Technical University, Ankara, TR-06800, Turkey. His research interests include statistical signal processing and its applications in array signal processing and radar/sonar signal processing.

References

[1] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals & Systems*, 2nd ed. Englewood Cliffs, NJ: Prentice Hall, 1999.

[2] A. V. Oppenheim, R. W. Schafer, and J. R. Buck, *Discrete-Time Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1999.

[3] A. Papoulis, *The Fourier Integral and Its Applications*. New York: McGraw-Hill, 1962.

[4] R. Bracewell, *The Fourier Transform & Its Applications*, 3rd ed. New York: McGraw-Hill, 1999.

[5] "Wikipedia contributors, Freeman Dyson," Wikipedia, 2020. Accessed: July 29, 2020. [Online]. Available: https://en.wikipedia.org/wiki/Freeman_ Dyson

[6] Y. Yamamoto, *From Vector Spaces to Function Spaces: Introduction to Functional Analysis with Applications*. Philadelphia: Society for Industrial and Applied Mathematics, 2012.

[7] M. J. Lighthill, *An Introduction to Fourier Analysis and Generalised Functions*. Cambridge, U.K.: Cambridge Univ. Press, 1952 (online republication 2012).

[8] D. G. Luenberger, *Optimization by Vector Space Methods*. Wiley-Interscience, 1969.

[9] G. Temple, "Theories and applications of generalized functions," *J. London Math. Soc.*, vol. s1-28, no. 2, pp. 134–148, 1953. doi: 10.1112/jlms/s1-28.2.134.

[10] G. Temple, "The theory of generalized functions," *Proc. R. Soc. Lond. A*, vol. 228, no. 1173, pp. 175–190, 1955. [Online]. Available: https://www.jstor.org/stable/99615

SP

DATES AHEAD

Please send calendar submissions to:
Dates Ahead, Att: Samantha Walter, Email: walter.samantha@ieee.org

Editor's Note

Due to changing situations around the world because of the novel coronavirus (COVID-19) outbreak, please double-check each conference's website for the latest news and updates.

©SHUTTERSTOCK/DMITRI KALVAN

The IEEE Statistical Signal Processing Workshop will be held 11–14 July in Rio de Janeiro, Brazil.

2021

MAY

International Conference on Information Processing in Sensor Networks (IPSN)
18–21 May, Nashville, Tennessee, United States.
General Chair: Ákos Lédeczi
URL: https://ipsn.acm.org/2021/

JUNE

IEEE Data Science and Learning Workshop (DSLW)
5–6 June, Toronto, Ontario, Canada.
General Chairs: Stark Draper and Z. Jane Wang
URL: https://conferences.ece.ubc.ca/dslw2021/#/

IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)
6–11 June, Toronto, Ontario, Canada.
General Chairs: Dimitri Androutsos, Kostas Plataniotis, and Xiao-Ping (Steven) Zhang
URL: https://2021.ieeeicassp.org/

Virtual: International Conference on Quality of Multimedia Experience (QoMEX)
14–17 June
General Chairs: Tiago H. Falk and Amy Reibman
URL: https://qomex2021.itec.aau.at/

Content-Based Multimedia Indexing (CBMI)
28–30 June, Lille, France.
Conference Chair: Chaabane Djeraba
URL: https://cbmi2021.univ-lille.fr

Digital Object Identifier 10.1109/MSP.2021.3054956
Date of current version: 28 April 2021

Picture Coding Symposium (PCS)
29 June–2 July, Bristol, U.K.
General Chair: David Bull
URL: https://pcs2021.org

JULY

IEEE International Conference on Multimedia and Expo (ICME)
5–9 July, Shenzhen, China.
General Chairs: Moncef Gabbouj, Houqiang Li, Guo-Jun Qi, and Yonghong Tian
URL: https://2021.ieeeicme.org/

IEEE Statistical Signal Processing Workshop (SSP)
11–14 July, Rio de Janeiro, Brazil.
General Chair: Rodrigo C. de Lamare
URL: http://ssp2020.cetuc.puc-rio.br

AUGUST

IEEE International Conference on Autonomous Systems (ICAS)
11–13 August, Montréal, Québec, Canada.
General Cochairs: Amir Asif and Arash Mohammadi
URL: https://2021.ieee-icas.org

SEPTEMBER

Sensor Signal Processing for Defence (SSPD)
14–15 September, Edinburgh, United Kingdom.
General Chairs: Mike Davies, Stephen McLaughlin, Jordi Barr, and Gary Heald
URL: https://sspd.eng.ed.ac.uk/

IEEE International Conference on Image Processing (ICIP)
19–22 September, Anchorage, Alaska, United States.
General Chair: Saif alZahir
URL: https://2021.ieeeicip.org

IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)
27–30 September, Lucca, Italy.
General Chair: Luca Sanguinetti
URL: https://www.spawc2021.com

SP

IEEE Machine Learning for Signal Processing Workshop 2021, 25 – 28 October, Gold Coast, Australia

General Chairs
- Abd-Krim Seghouane
University of Melbourne
- Mohammed Bennamoun
University of Western Australia
- Jonathan Manton
University of Melbourne

Program Chairs
- Dong Xu
University of Sydney
- Liang Zheng
Australian National University
- Wen Li
University of Electronic Science and Technology Chine

Plenary Chairs
- Ba-ngu Vo
Curtin University
- Tongliang Liu
University of Sydney

Tutorials Chairs
- Hamid Laga
Murdoch University
- Qian Yu
Beihang University

Special sessions Chair
- Lu Sheng
Beihang University

Finance Chair
- Luping Zhou
University of Sydney

Student Prize Chair
- Chunhua Shen
University of Adelaide

Data Competition Chairs
- Ercan Kuruoglu
Consiglio Nazionale delle Ricerche
- Danilo Comminiello
Sapienza University of Rome

Publicity Chair
- Navid Shokouhi
Global Kinetics

Publication Chairs
- Guo Lu
Beijing Institute of Tectnology
- Jing Zhang
Beihang University

Advisory Committee
- Zheng-Hua Tan
Aalborg University
- Murat Akcakaya
University of Pittsburgh
- Bhaskar Rao
University of California San Diego
- Raviv Raich
Oregon State University

The 31st MLSP workshop in the series of workshops organized by the IEEE Signal Processing Society MLSP Technical Committee will take place in Gold Coast, Australia. If conditions allow, we plan to have a conference running in a hybrid format with a virtual program for attendees that can't attend the conference and in person attendance for attendees that can attend. The conference will present the most recent and exciting advances in machine learning for signal processing through keynote talks, tutorials, special and regular single-track sessions as well as matchmaking events. The presented papers will be published in and indexed by IEEE Xplore.

All submitted papers are reviewed by experts and only a proportion is accepted maintaining a high quality scientific meeting. The scope of the workshop includes basic theory, methods and algorithms, and applications in the following areas:

Theoretical and application Topics:
- Learning theory and algorithms
- Information-theoretic learning
- Deep learning techniques
- Distributed/Federated learning
- Dictionary learning
- Graphical and kernel methods
- Learning from multimodal data
- Independent component analysis
- Matrix factorizations/completion
- Reinforcement learning
- Transfer learning
- Source separation
- Reinforcement learning
- Subspace and manifold learning
- Sequential learning
- Self-supervised and semi-supervised learning
- Tensor-based signal processing
- Sparsity-aware processing
- Pattern recognition and classification
- Music and audio processing
- Applications of machine learning

Submission of papers: Prospective authors are invited to submit 6 pages full-length papers, including figures and references. Papers are submitted electronically at 2021.ieeemlsp.org

Special session proposals: Special session proposals are submitted through the 2021.ieeemlsp.org. They must include a topical title, rational, session outline, contact information and the list of invited papers. Special session authors are referred to the workshop website for additional information regarding submissions.

Important Dates:

- Deadline for special session proposals submission — April 30th, 2021
- Special sessions selected — May 07th, 2021
- Deadline for 6-page paper submissions — May 31st, 2021
- Notification of paper acceptance — July 31st, 2021
- Camera ready upload — August 31st, 2021
- Start of conference — October 25th, 2021

Venue: The Gold Coast is a coastal city located in the South East of Queensland, Australia. The city is 94 km (58 mi) south of the state capital Brisbane. With more than 300 days of sunshine each year allowing you to explore more than 70 km of unspoiled coastline and beaches as well as 100,000 hectares of world heritage rainforest; Gold Coast is today a leading touristic destination. It enjoys 30 different beaches each with their own personality, from world-famous surf breaks and cosy coastal settings in the south to family fun swimming spots and surf-life saving action in the north. At your fingertips is every on-water activity imaginable! With its sunny subtropical climate, surfing beaches, canal and waterway systems, its high-rise dominated skyline, theme parks, nightlife and rainforest hinterland, Gold Coast will turn your stay into lifetime memories.

Digital Object Identifier 10.1109/MSP.2021.3070727

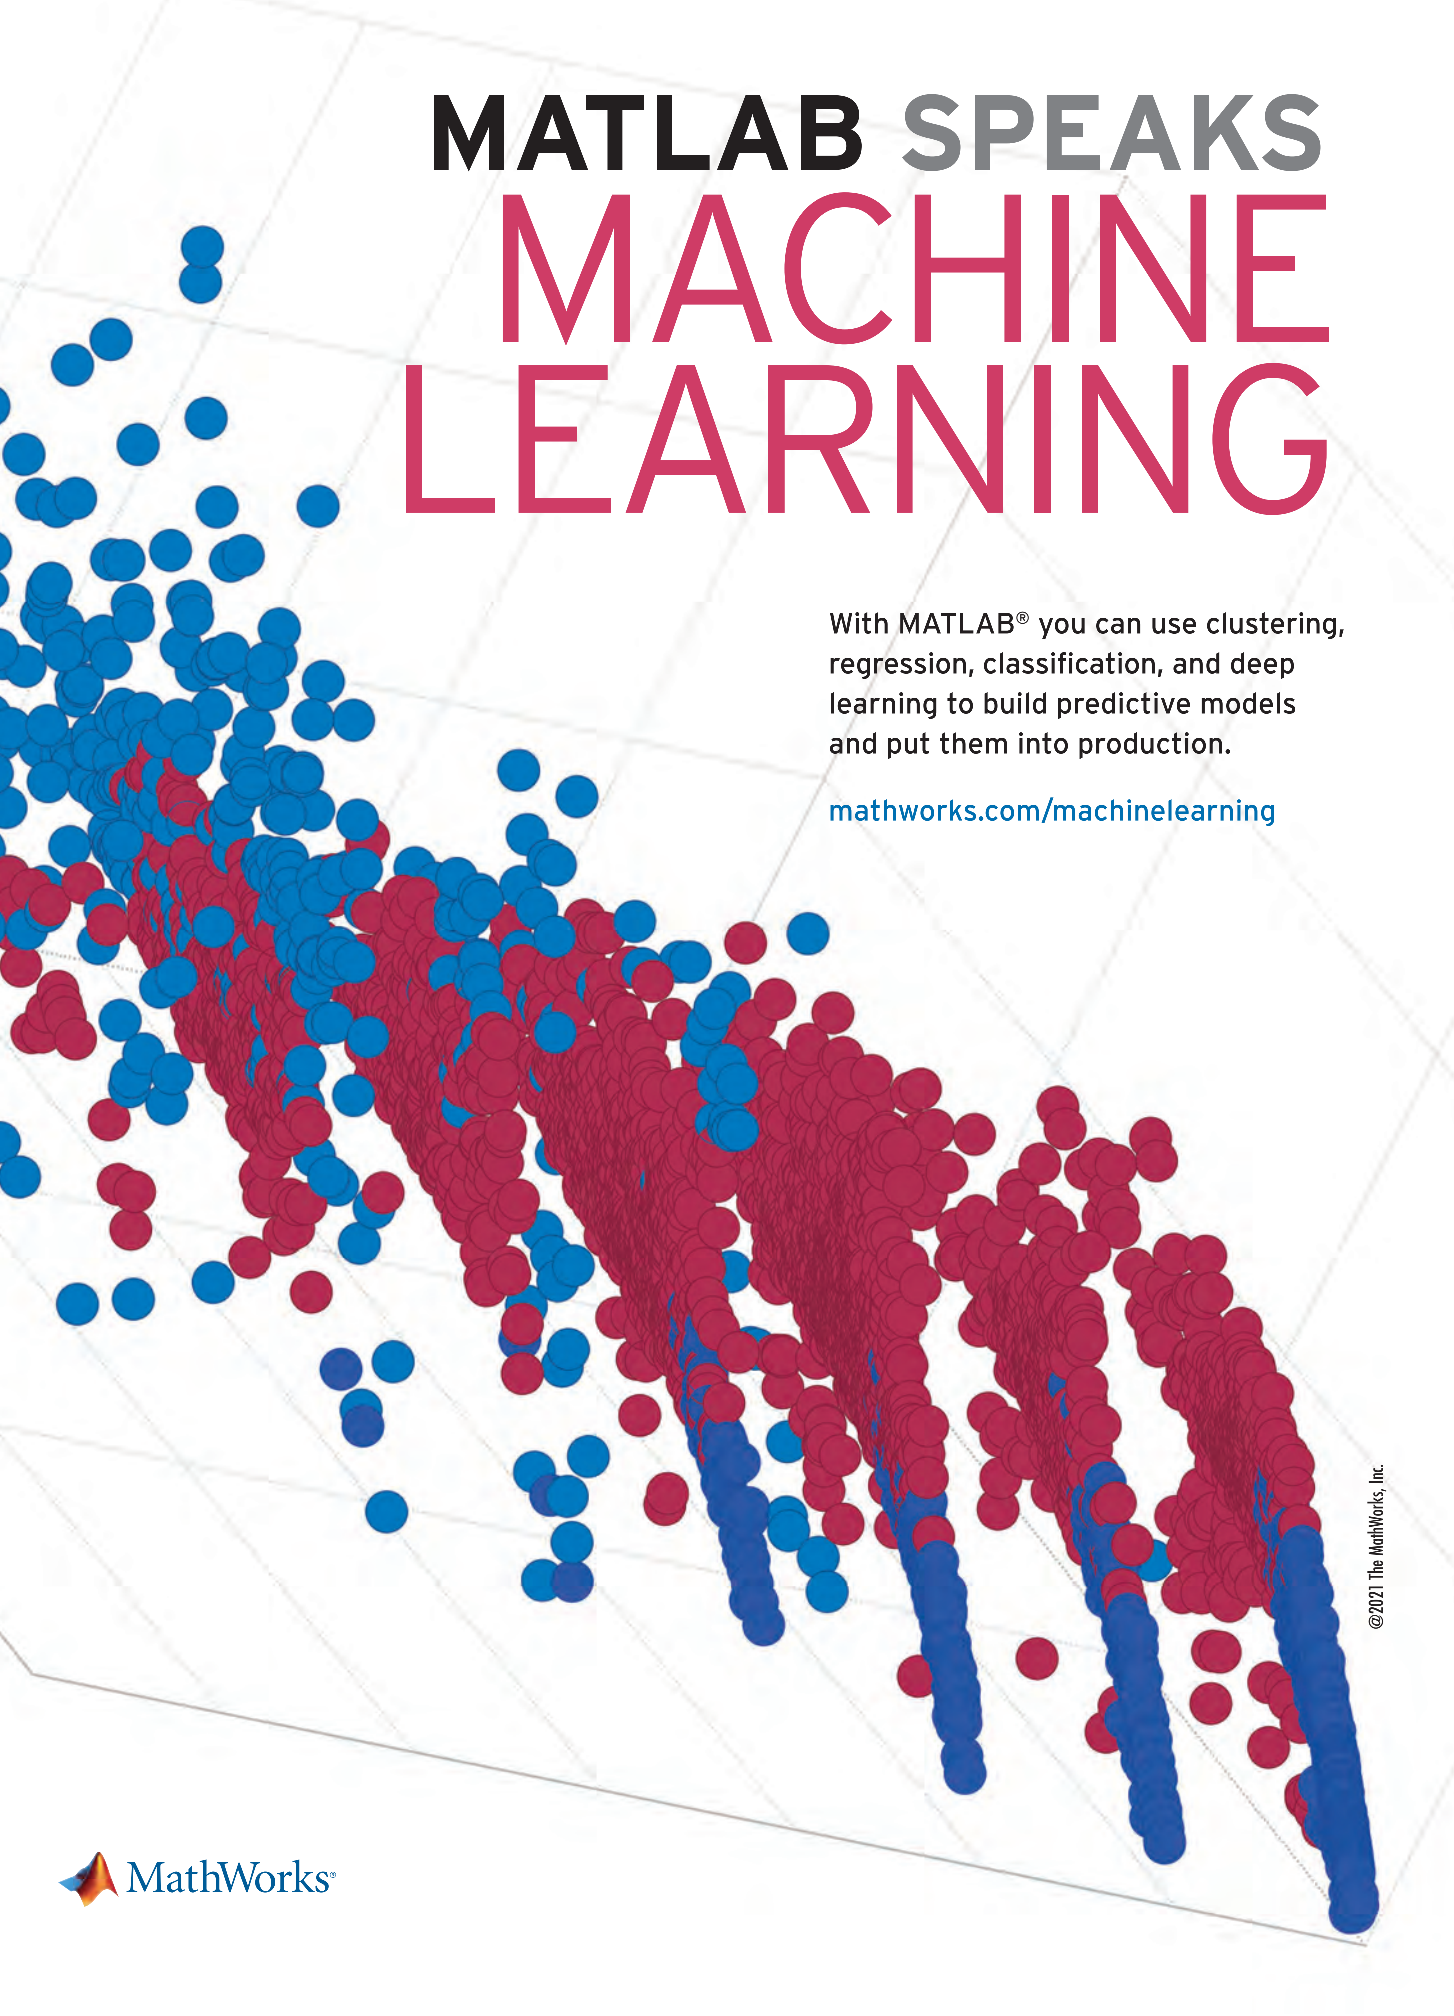
MATLAB SPEAKS
MACHINE
LEARNING
With MATLAB® you can use clustering, regression, classification, and deep learning to build predictive models and put them into production.
mathworks.com/machinelearning
@2021 The MathWorks, Inc.
MathWorks®