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

Balancing Artificial and Natural Intelligence

For me and, probably, many readers, each issue of *IEEE Signal Processing Magazine* (*SPM*) is the opportunity and pleasure to learn something new in the area of signal and image processing. In addition to lecture notes, tips-and-tricks articles, special reports, and so on, which propose interesting and clever solutions to typical signal or image processing problems, the feature articles and special issue provide tutorial-like articles on various mature or fast-developing domains.

In this *SPM* March issue, if you are not an expert in advanced communications, you will discover the concept of a reconfigurable intelligent surface (RIS). This device is based on an ultra-fine material organized in an array of reconfigurable elements, whose scattering, absorption, reflection, and diffraction properties can be changed with time and controlled by software. Thus, an RIS acts like a smart antenna able to focus transmitter signals directly toward different receivers, thus optimizing the signal/noise ratio and enhancing the quality of the communication.

SPM's March issue includes seven special issue articles focused on "Deep Learning in Biological Image and Signal Processing," which is presented in more detail in the "From the Guest Editors" column (page 24). In fact, artificial intelligence, machine learning, and especially deep learning (DL) are ubiquitous in all data science applications; the biological and biomedical domains are no exception to this, and for good reason. DL algorithms have, in recent years, proven their ability to solve many problems, like classification, segmentation, feature extraction, prediction, and so on with outstanding performance, provided that you have both a very large database with good-quality data and high computational power.

In the medical domain, these requests have been satisfied by worldwide collaborations providing huge databases obtained by sharing data from among hospitals located in many countries and, also, access to shared superpowered calculators. For instance, we now have very huge brain imaging databases coming from patients of different ages, healthy or not, and we are able to process them, thus opening a new understanding of the brain. As explained in "Explainable Artificial Intelligence for Magnetic Resonance Imaging Aging Brainprints" (page 99), the prediction of the brain's age with DL allows the early detection of neurodegenerative disease, accurate tracking of the disease progression, and personalization of the therapy to each patient.

The first methods in DL were based on supervised learning, which needs to have annotated data: this is a limitation for medical data, where there is basically no ground truth, as attested to by discrepancies among the annotations of different experts! For overcoming this issue, other machine learning methods, unsupervised or self-supervised, can then be used successfully, as detailed in "Unsupervised Deep Learning Methods for Biological Image Reconstruction and Enhancement: An Overview From a Signal Processing Perspective" (page 28).

A DL algorithm is basically a black box characterized by millions of parameters that must be tuned by learning. This implies at least two essential problems, especially in medical applications, which are widely addressed and discussed in all of the articles: the (1) robustness of the results (i.e., weak sensitivity of the outputs to small changes in the data), which impacts their reliability and reproducibility, and (ii) explainability and interpretability, which are fundamental for the physician to be confident in the results, understand them, and be able to explain them to the patient. In the future, we can hope that explainability could be a way to gain insights about a disease and mechanisms for disorders so that one can come up with preventive measures and new treatment strategies.

> **An RIS acts like a smart antenna able to focus transmitter signals directly toward different receivers, thus optimizing the signal/noise ratio and enhancing the quality of the communication.**

International (1994–2011) and was affiliated with the International Computer Science Institute in Berkeley, California—most recently as an external fellow. He received a Diplom (master's) degree from the Technical University of Munich (1988) and a Ph.D. in computer science from UC Berkeley (1994) for his thesis work on probabilistic parsing and grammar induction.

Dr. Stolcke served as an associate editor of *IEEE Transactions on Audio, Speech, and Language Processing* (2000–2002), a coeditor of *Computer Speech and Language* (2003–2006), and an editorial board member of *Computational Linguistics* (1997–1999). He has organized special sessions and workshops at ICASSP and Association for Computational Linguistics conferences. He served on the SPS Speech and Language Processing Technical Committee (2013–2019) and is a Fellow of both IEEE (2011) and the ISCA (2013).

Dr. Stolcke has made contributions to machine learning and algorithms for speech and language processing, including conversational speech recognition, speaker recognition and diarization, and paralinguistic modeling. He developed the entropy-based pruning method for N-gram language models and designed and open-sourced the widely used SRI Language Modeling toolkit. He pioneered several methods for using automatic speaker recognition by-products for speaker recognition and conceived the DOVER algorithm for combining multiple diarization hypotheses.

Dr. Stolcke's current work is focused on exploiting the full range of speech communication in speech and speaker understanding and making conversational speech agents more natural and contextually aware. Dr. Stolcke's lecture topics include speech recognition and understanding for conversations and meetings, speech technology for advanced conversational agents, and recent advances in speaker recognition and diarization.

SP

FROM THE EDITOR

Finally, in "A Practical Guide to Supervised Deep Learning for Bioimage Analysis" (page 73), I noticed that a very sound question is addressed, even if it is seen as iconoclast in the deep learning culture: is DL actually needed? The authors recommend to start by exploring classical (simpler, faster, and lower-cost computation, power, and memory) and non-DL machine learning methods before considering DL.

Finally, I would like to finish with one general observation concerning robustness and reproducibility. This observation is sounded in biomedical signal and image processing but also far beyond them. In many application domains, current data processing solutions are more and more complex, characterized by a pipeline of many processing blocks, where each block is usually tuned by one or many more (millions in DL!) parameters. Of course, processing the same data with different pipelines can provide completely different results! Much more, in the open science framework, it is usual to use the same pipeline as other teams, but all of its parameters are not always accurately defined and shared; in complex pipelines, changing a few of them, even just one in a unique block (e.g., a threshold in a preprocessing block), can then provide very different results [1]. This issue is critical, especially for international collaborations among scientists who want to share their data, methods, and results or simply do relevant benchmarks.

Do not let artificial intelligence blindly drive our works. Let's allow natural intelligence to guide our thinking and our solutions.

Reference
[1] A. Bowring, T. Nichols, and C. Maumet, "Isolating the sources of pipeline-variability in group-level task-fMRI results," *Hum. Brain Mapping*, pp. 1–17, Nov. 13, 2021, doi: 10.1002/hbm.25713. **SP**