Biomedical Imaging and Analysis in the Age of Big Data and Deep Learning

By **JAMES S. DUNCAN**, *Life Fellow IEEE*
Guest Editor

MICHAEL F. INSANA, *Life Fellow IEEE*
Guest Editor

NICHOLAS AYACHE, *Member IEEE*
Guest Editor

Imaging of the human body using a number of different modalities has revolutionized the field of medicine over the past several decades and continues to grow at a rapid pace [2]. More than ever, previously unknown information about biology and disease is being unveiled at a range of spatiotemporal scales. Although results and clinical adoption of strategies related to the computational and quantitative analysis of the images have lagged behind development of image acquisition approaches, there has been a noticeable increase of effort and interest in these areas in recent years [6]. This special issue aims to define and highlight some of the "hot" newer ideas that are in biomedical imaging and analysis, intending to shine a light on where the field might move in the next several decades, and focuses on emphasizing where electrical engineers have been involved and could potentially have the most impact. These areas include image acquisition physics, image/signal processing, and image analysis, including pattern recognition and machine learning. This issue focuses on two themes common in much of this effort: first, engineers and computer scientists have found that the information contained in medical images, when viewed through image-based vector spaces, is generally quite sparse. This observation has been transformative in many ways and is quite pervasive in the articles we include here. Second, medical imaging is one of the largest producers of "big data," and data-driven machine-learning techniques (e.g., deep learning) are gaining significant attention because of improved performance over previous approaches. Thus, data-driven techniques, e.g., formation via image reconstruction [11] and image analysis via deep learning [8], [9], are gaining momentum in their development.

This special issue aims to define and highlight key ideas in biomedical imaging and analysis that shed a light on where the field is headed in the next decade.

The set of articles included examine the capability of image science to explore the complexity of life systems, from bacterial colonies to human medicine. This goal has challenged biological and medical imaging researchers to develop sensing techniques capable of tracking cellular communications over a large range of spatiotemporal scales to explore the hierarchy of properties emerging from complex living systems. The search for deeper understanding and clearer diagnostic assessments is driving technology into higher dimensional spaces. Ideas that began with multimodality approaches for imaging and treating cancer and cardiovascular disease have expanded into developing techniques that reveal the systemic role of the microbiome in healthy cells and disease cells, the topology of brain connectivity and biochemistry in cognition, and cognitive computing in image formation and interpretation where human pattern recognition and model-based image formation methods are hitting their limits. The limitations encountered when modeling instruments

Digital Object Identifier 10.1109/JPROC.2019.2956422

as linear systems can be overcome using data-driven approaches now offered through a range of machine-learning techniques. Yet many sense that the most useful and robust models may involve some mixture of model-based and data-driven approaches. The articles that are included focus on a collection of topics, which we feel are important areas for the future of biomedical imaging. These are spread across ten contributions from ten different sets of authors that are detailed below. In this "Scanning the Issue" article, we first try to set the stage for the Special Issue by briefly reviewing the recent history of the terminology used in the fields of big data, machine learning, and deep learning in the context of medical imaging. Then, we will introduce and summarize the ten articles by grouping them into categories of: A) modality-centric image acquisition efforts, including image reconstruction and B) efforts that are more focused on image analysis and image-guided intervention. We will conclude by summarizing some of the cross-cutting themes of the contributions.

I. BACKGROUND AND TERMINOLOGY: PATTERN RECOGNITION, BIG DATA, MACHINE LEARNING, AND ARTIFICIAL INTELLIGENCE

Although the history of radiology began with Wilhelm Roentgen's taking the first X-ray image (of his wife) in 1895 and evolved through 1913 with the invention of mammography and with the first cerebral angiogram taken in 1927, modern-day medical imaging began to take shape in the 1950s with the invention of positron emission tomography (PET) and ultrasound imaging. Perhaps the key defining moments of computational medical imaging came with the 1970s as Godfrey Hounsfield (an Electrical Engineer) and Allan Cormack invented the first computed tomography (CT) scanner in 1972 and then the first commercial magnetic resonance imaging (MRI) scanner was developed by Raymond Damadian in 1977. Somewhat in synchrony with the medical imaging equipment developments in the 1970s and beyond, and the age of digital computers, were the development of the general techniques and terms of digital image processing and pattern recognition. Digital picture and image processing were developed in the 1960s, mostly at places like Caltech/JPL, MIT, and Bell Laboratories, and were most often associated with imaging and exploration of outer space. Interestingly, in 1994, "Digital Image Processing—Medical Applications" was an "Inducted Technology" into the *Space Technology Hall of Fame* [10], illustrating at least one connection between these fields. But perhaps most germane to the current article is the evolution of the term "pattern recognition" and its relationship with the notions of machine learning and deep learning, all ideas that influence the ten articles are included in this Special Issue. Pattern recognition first became an area of study in the early 1970s, perhaps best exemplified by the classic textbook *Pattern Classification* by Duda and Hart [7], which was first published in 1973. This textbook and this very field most definitely evolved out of the general field of electrical engineering (EE) as a type of intelligent (digital) signal processing, thereby developing into a subfield solidly based in the EE university curricula of this time frame. The early goals of work in pattern recognition were to develop algorithms that would be implemented in software or hardware to perform intelligent tasks similar to what a human could perform, for example, picking out trends in an electrocardiogram or finding objects in an image. Decision making was typically done using features extracted from the data and run through heuristic or logical decision trees or discriminant analyses based on Bayesian statistical methods. In the 1980s and 1990s, the field gained more attention from computer scientists and decision-making algorithms moved toward using increasing amounts of data, with the goal of less human intervention being required, leading to the coining of the term "machine learning." One of these strategies was based on multiple layers of simple decision-making nodes that loosely tried to mimic human brain networks, known as computational neural networks. It is most interesting to see the evolution of the use of these terms during the early years and beyond the 21st century, with pattern recognition becoming less and less popular and machine learning and deep learning becoming ubiquitous as evident from Fig. 1 (see also the earlier comparison and discussion in [5]).

Certainly related to the above are the ideas of "Big Data" and "AI," the latter of which has seen its use go up and down over the last half-century but (in at least some definitions) has been used to encompass everything and anything related to the idea of computers and learning algorithms.

II. OVERVIEW OF THE SPECIAL ISSUE

As discussed above, this issue contains ten contributions from some of the most active investigators in the medical imaging field who are using computational strategies to affect their approaches and to improve the utility of information contained within, and derived from, medical images. The topics have been addressed by a group of authors who are mostly from separate institutions or companies and come from a number of regions across the globe. The articles are all driven by applications but show a range of technology development. We feel that these can mostly be clustered into two large subcategories that make up Sections II-A and II-B: A) work related to image acquisition and formation and B) work related to image analysis and image-guided intervention, including the integration of nonimage data such as genomics.

A. Image Acquisition and Formation

In this section, five different sets of contributors look at how data-driven/

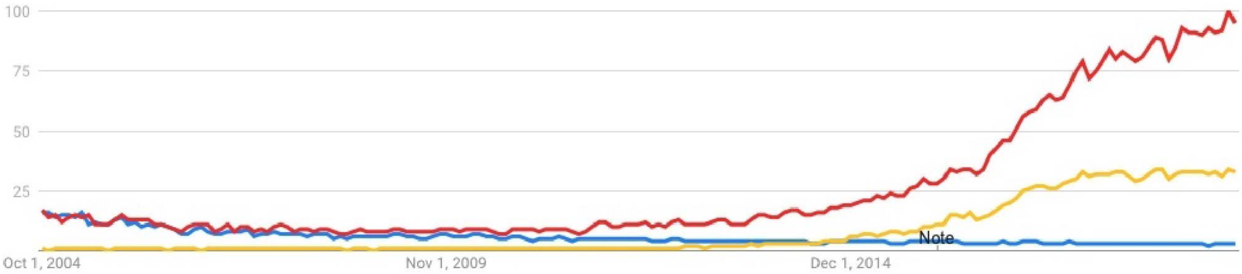

Fig. 1. ***Search of terminology use trends from Google Trends (https://trends.google.com/trends/): pattern recognition (blue) versus machine learning (red) versus deep learning (yellow) between 2004 and 2018. These scores awarded on this "interest over time" line graph express the popularity of that term over a specified time range. Google Trends scores are based on the absolute search volume for a term, relative to the number of searches received by Google. The scores have no direct quantitative meaning. For example, two different terms could achieve scores of 100 in the same month, but one received 1000 search requests, while the other received 1 000 000. This is because the scores have been scaled between 0 and 100. A score of 100 always represents the highest relative search volume. These monthly scores are calculated on the basis of the average relative daily search volume within the month. A rising or declining line does not necessarily indicate a change in the popularity but rather likely indicates that general search use has increased or decreased over the time range.***

machine-/deep-learning affects the formation of images.

First, in the article "Deep learning in ultrasound imaging," van Sloun *et al.* explore the use of deep, data-driven learning on all aspects of ultrasound imaging, ranging from ideas that are at the interface of raw signal acquisition (including adaptive beamforming) and image formation to learning compressive codes for color Doppler acquisition and to learning strategies for performing clutter suppression. They offer a provocative vision of the future of ultrasound based on extremely portable and intelligent imaging that facilitates smart, wireless probes that are aimed at a range of very specific applications. The authors first note that sonographic instruments are limited by the high volume of data that must be processed to implement the available methods. For example, implementation of high-frame-rate 3-D sonography with advanced beamforming and sensitive blood- and tissue-motion imaging capabilities has placed extraordinary demands on receive beamforming and related signal processing. In addition, the use of oversampling to provide fine-scale phase adjustments during dynamic-receive focusing limits the frame rate and/or depth of tissue penetration. The authors describe new methods involving echo modeling where copies of the sound pulse with unknown amplitudes and delays sparsely represent the response of the system to tissue scatterers. This sparse-data model can lead to significantly reduced sampling rates without limiting performance. Compressed sensing techniques also eliminate the need for analog microbeamforming techniques now applied to speed processing at the cost of image quality. They show how efficient sampling enables the addition of new computational imaging technologies throughout many applications of medical sonography. In addition, they also show a variety of examples in their article, concluding that an integration of model- and data-driven approaches are often the best. More specifically, highlights from this article first include the idea that front-end beamforming strategies could be designed using deep learning by learning the delays and the apodizations by creating dedicated delay layers. Furthermore, stacked autoencoders or convolutional neural nets could be used to generally map pre-delayed channel data to beamformed outputs. Similar architectures could learn how to take raw RF data into optimized B-mode images. Deep networks could also be used for estimating spectra in spectral Doppler applications and some of this work has been done already. As in the optical microscopy applications described below, deep learning and sparsity can be used to develop super-resolution ultrasound. In clinical echocardiography, recognizing optimal views has been achieved using deep networks.

In the second article, titled "Deep learning-based image reconstruction and enhancement in optical microscopy," de Haan *et al.* develop an overview of efforts to advance the field of computational microscopy and optical sensing systems for microscopy using deep neural networks. They first overview the basics of inverse problems in optical microscopy and then outline how deep learning can be a framework for solving these problems, typically through supervised methods. Then, they focus on the use of deep learning to try to obtain single image super resolution and image enhancement in these data sets. Here, they describe work that is able to extrapolate missing spatial frequencies and obtain an extended depth of field (DOF), noting how these strategies employ different deep neural network architectures, including generative adversarial networks (GANs). Perhaps one of the most exciting areas noted in this article is the new use of deep learning to address the reconstruction of single molecular localization images at extremely high spatial resolution, known as photoactivated localization microscopy (PALM) and stochastic optical reconstruction microscopy (STORM). In both cases, when deep learning was employed, the techniques performed significantly faster while maintaining the same high spatial resolution. As the authors note, microscopy is an ideal field to apply deep learning, since the experimental data sets are typically obtained under very controlled conditions, including fixed and repeatable illumination and focus. The authors conclude that new smart systems are possible that could be engineered to solve specific image analysis tasks by integrating the acquisition system with the analysis algorithm that could even predict which measurement is required next. Such task-specific, personalized, "thinking" imaging systems are ideas that are in common with the conclusions of other articles in this issue.

The third article in this group, "Machine learning in PET: From photon detection to quantitative image reconstruction," by Gong *et al.* addresses the uses of machine-learning strategies in the area of PET imaging. Here, the authors discuss applications of machine learning to PET, PET-CT, and PET-MRI multimodal imaging. They describe the impact of machine learning both at the detector stage and for quantitative image reconstruction. Given that there are roughly equal contributions from true coincident counts, scattered photon events, and random noise events in PET detection, and considering the challenges imposed by low intrinsic detection efficiencies, effective application of signal processing to the detector signals is essential in the search for an optimal balance among patient dose, scan time, and image quality. Detector processing involves determining the timing and position of absorption events within each detector crystal based on the distribution of scintillation photons recorded by photodetectors. Given the many influences on light distribution within the detector, a broad range of machine-learning approaches from traditional statistical pattern recognition algorithms to convolutional neural networks (CNNs) have shown much promise at improving sensitivity. As the authors note in their abstract, fast waveform digitizers are now available, and machine learning has been used to actually estimate the position and arrival time of high energy photons. They also discuss how a broad array of statistical methods and neural network applications is improving performance of attenuation and scatter correction algorithms, as well as integrating patient priors into reconstructions based on a constrained maximum likelihood estimator. In the reconstruction portions of the image-formation algorithms, machine and data-driven learning has been used to correct for scatter and attenuation while reducing noise. As indicated in reference 166 of their article, new ideas in the field moving forward include trying to recognize pattern relationships between low and high count data in order to estimate one day high count data from limited count data.

MRI has become ubiquitous as a go-to modality in many areas of medical imaging fields. Rather than overviewing this entire field, we chose to ask the authors of the fourth article, in this section, to focus on one interesting area, which resulted in a contribution titled "Machine learning for rapid magnetic resonance fingerprinting tissue property quantification" by Hamilton and Seiberlich. In this approach, a single rapid MRI acquisition can produce quantitative maps of multiple tissue properties simultaneously. The magnetic resonance fingerprinting (MRF) approach was initially developed with notions of compressed sensing and sparsity in mind to generate a dictionary of signals using Bloch equation simulations. However, the latest techniques described in this contribution describe how machine learning can now accelerate the extraction of quantitative maps from the MRF results. More specifically, the authors describe how neural networks may accelerate dictionary generation, which is crucial for applications that quantify many tissue properties simultaneously or require frequent calculation of new dictionaries. In addition, they note how machine learning may permit faster, more robust pattern recognition by bypassing dictionary generation and directly estimate tissue property values from the measured data. In fact, a version of the same techniques may provide faster and more robust reconstructions of tissue property maps that could aid the clinical translation and adoption of MRF. These ideas support the integration of acquisition and analysis theme noted in the previous two paragraphs.

The fifth and final contribution in this section, titled "Image reconstruction: From sparsity to data-adaptive methods and machine learning," by Ravishankar *et al.*, overviews how sparsity, data-driven methods, and machine learning have, and will continue to, influence the general area of image reconstruction, cutting across modalities. In this mathematically solid article, the authors review the basic strategies currently in use as well as describe the work now aimed at trying to interpret what the deep-learning models are actually doing. The authors describe the strong influence that sparse and reduced-rank data models have made on model-based and data-driven image reconstruction techniques. By summarizing the history of image reconstruction, they show how image quality is driving the evolution of ideas toward the cutting-edge methods available today. The greatest impact of sparsity on MR is to reduce scan times, while that on CT is to reduce patient dose. Referring back to the ultrasound discussion above, we notice that both the article by van Sloun *et al.* as well as this fifth contribution by Ravishankar *et al.* describe how sparse and low-rank models can simultaneously recover mixed signal components through separate constraints. Both articles describe structured low-rank models, e.g., the L + S decomposition, and the associated algorithms now in use. Source separation is valuable for eliminating unwanted motion artifacts, for example, or separating anatomical from functional signals. Furthermore, we note that Ravishankar *et al.* describe the current ideas on the use of deep-learning strategies for general image reconstruction. One particularly interesting class of popular hybrid domain approaches is based on the CNN penalty and plug-and-play model. Finally, related to the previous articles in this section, the authors note that they expect that the next generation imaging systems would leverage learning in all aspects of the imaging system, learning to optimize the underlying models for efficient and effective reconstruction and analytics, including ideas of classification, segmentation, and disease feature detection. Such a design would permit the data acquisition, reconstruction, and analytics components to be done jointly in an end-to-end manner to maximize performance in specific clinical tasks and allowing for both radiologist and patient inputs in the learning process.

B. Image Analysis, Including Integration of Non-Image (e.g., Genomics) Data and Image-Guided Therapy/Intervention

In this section, another five sets of contributors look at how data-driven ideas have affected a variety of image analysis and image-guided intervention areas.

The first article titled "Model-based and data-driven strategies in medical image computing," by Rueckert and Schnabel, is a most interesting and informative historical overview as to how the medical image analysis and image computing field has developed, starting with model-based approaches and then evolving to today's current emphasis on data-driven/deep-learning-based efforts. The authors do an excellent job comparing these concepts and highlighting the advantages and disadvantages of each style of work. Most elucidating, given current efforts in the field, is the authors' insight that while data-driven, deep-learning approaches often can outperform the more traditional model-based ideas, the notion of using these techniques in clinical scenarios has led to a number of challenges, including the ideas that the approaches may be brittle/nonrobust to new examples and data are not easily generalizable and are almost impossible to interpret or explain. Table 1 of this article very clearly lays out comparisons between traditional, model-based, and current data-driven techniques and the article ends with provocative thoughts about how a new diagnostic, data-driven pipeline could be created that goes directly from images to diagnostic recommendations. When compared with previous visions of the future in medical image analysis, it may be of interest for the reader to look back at the article by Duncan and Ayache [1] published in the IEEE Transactions on Pattern Analysis and Machine Intelligence in 2000.

The second article in this section, "Brain imaging genomics: Integrated analysis and machine learning," by Shen and Thompson, describes applications of novel and traditional data-science methods to the study of "brain imaging genomics." One could see this as a further development of image-derived-only quantitative analysis. In the work noted here, the authors talk about how researchers combine diverse types of high-volume data sets, which include multimodal and longitudinal neuroimaging data and high-throughput genomic data with clinical information and patient history, to develop a phenotypic and environmental basis for predicting human brain function and behavior. They examine three categories of machine learning for brain imaging genomics: heritability of traits, learning the imaging-genomic associations, and applying this information for predicting behavior. This work assembles multivariate statistical and network-based techniques that enable the authors to work within huge arrays of imaging, omics, and medical databases as necessary to overcome formidable statistical and computational challenges. Their tools include four categories of regularized regression analysis, where weight matrices are adjusted to minimize a loss function comparing imaging and genomic data when identifying functional associations. The loss function is regularized not only to minimize overfitting, as usual, but also to reduce data dimensionality for irrelevant associations. Sparse and low-rank models populate the high-dimensional weight matrices. Dimensionality reduction is essential to increase the statistical power of predictions given modest sample sizes.

The third article related to image analysis topics discusses work in an area that has been studied for many years—computer-aided diagnosis (CAD) of breast cancer—but now looking at the problem through the lens of machine/deep learning. In "Comparison of breast MRI tumor classification using human-engineered radiomics, transfer learning from deep CNNs, and fusion methods," Whitney *et al.* have further taken a unique look at recent approaches, especially focusing on the capabilities of deep neural networks to perform transfer learning. In this manner, features that have been derived and found in the network layers of architectures intended to find useful features in natural image classification tasks are then used for classification tasks related to find benign and malignant lesions from breast images. CNNs have been designed to perform exactly such tasks and are discussed and overviewed in detail. The primary goal of the work described was to comprehensively build on the prior research of the authors, who have long worked in this area, and evaluate the classification performance of different classifiers, in the task of classification of lesions as benign or malignant, constructed using human-engineered radiomic features and two variations on transfer learning: features extracted from pretrained neural networks and features extracted after fine-tuning a neural network for a specific classification task. Four different associated fusion classifiers were evaluated, all using dynamic contrast enhanced (DCE) MRI data of the human breasts.

A fourth article in this section, "Wireless capsule endoscopy: A new tool for cancer screening in the colon with deep-learning-based polyp recognition," by Jia *et al.*, integrates notions of image acquisition and image analysis for use in cancer screening through wireless capsule endoscopy (WCE). Here, machine- and deep-learning approaches are being developed to assist in automated polyp recognition/detection and analysis that will enhance diagnostic accuracy and efficiency of this procedure that is a critical tool for use in the clinic. WCE allows direct visualization of the entire colon without patient discomfort but manual review is tedious and labor-intensive. Computational methods for automated analysis of polyps have great potential here and the use of deep-learning methods for these tasks has gained traction and promises better accuracy and computational speed in the years ahead. The majority of these approaches employ CNNs, but a recent work includes autoencoders and GANs for finding polyp bounding boxes as well as transferring features learned from natural images for use in tumor recognition and classification studies. In Table 1 of this article, the authors perform an excellent job of summarizing the range of deep-learning methods applied to WCE polyp recognition that can serve as a guide to work in this subfield.

In the fifth and final article of this section, titled "CAI4CAI: The rise of contextual artificial intelligence in computer-assisted interventions," Vercauteren *et al.* take a close look at the use and the rise of Contextual artificial intelligence (AI) in Computer Assisted Interventions in the medical imaging field. The primary challenges in this subfield include how to incorporate the range of prior knowledge and instantaneous sensory information from experts, sensors, and actuators, as well as learning how to develop a representation of the surgery or intervention among a mixed human–AI team of actors. In addition, the authors describe how to design interventional systems and associated cognitive shared control schemes for online uncertainty-awareness while making decisions in the operating room (OR) or in the interventional radiology (IR) suite, tasks that are critical for producing precise and reliable interventions. Much of all of this involves the integration of all sorts of medical data, including images for guidance of the interventions or surgeries, and has led to the coining of the term "surgical data science."

III. CROSS-CUTTING THEMES

As we look across the ten contributions to this Special Issue, we notice several solid themes that have emerged. First, in most articles, the authors make an important observation that data-driven, deep learning may likely unify the typically modularized design of imaging systems. The sensing, acquisition/reconstruction/formation, and analysis steps that permit medical imaging to be used to quantify disease may now very well become more integrated, and we may also begin to see complete end-to-end system designs that are made more task-specific to optimize performance. We saw this clearly in the above discussion in Section III-A and in the articles related to PET, optical, and MRI systems, as well as in most all of the articles in Section III-B. Along with this, we note that several of the contributions, notably Hamilton and Seiberlich, Whitney *et al.*, and Ravishankar *et al.* in Section III-A, as well as Rueckert and Schnabel and Vercauteren *et al.* in Section III-B have begun to consider the integration of ideas and designs based on more traditional (and sometimes sparse) model-based methods with data-driven deep-learning methods for a variety of reasons, including as an approach to include helpful priors and to handle problems where limited amount of training data are available. The editors even wonder whether this is a current trend in the last year or so that helps explain the leveling out of the frequency of use of the term "deep learning" in 2019 (the yellow curve) in the Google Trends graph shown in Fig. 1.

Second, we also note that many of the articles in the issue make use of sparse methods. This is particularly the case with respect to sensing and acquisition, as compressive sensing (CS) makes it possible to acquire data from patients at rates far below the Nyquist limit without incurring a significant loss of information. Basically, all of the articles in Section III-A consider sparsity in one form or another. One could further observe that if the acquired data are sparse in some domain and the acquisition is incoherent as viewed through an isometric property, there are nonlinear optimization methods capable of fully recovering the patient features and measurements that one would seek from the images [3], [4]. This advance has given new life to the pursuit of fast-acquisition MR, low-dose CT, 3-D ultrasound with 2-D arrays, and all measurements fundamentally limited by the electronic data transfer and processing of high-volume data sets. The search for sparse representations of data has benefited from an expanded view of standard decompositions like singular value decomposition (SVD), so that we now reach beyond basis sets to include frames and dictionaries [12]. The latter two representations loosen the rules of decomposition so we can accept a richer palette for very sparse representation. With sparse data, we also achieve greater source and/or class separability. Several of the articles in Section III-B make this specific point with respect to image analysis (especially the article by Rueckert and Schnabel) and a number of articles in both Sections III-A and III-B discuss these points and show how efficient data acquisition aids in statistical classification and training neural networks. As we learn to recognize the intrinsic dimensionality of medical conditions, we will improve the efficiency of our healthcare system while lowering costs and improving patient outcomes. This is the true value of the Big Data revolution in medical imaging.

Acknowledgments

The guest editors would like to thank each one of the invited authors for preparing their most valuable

contributions and each one of the reviewers for their efforts during the review process. They would also like to thank the Editorial Board of the PROCEEDINGS OF THE IEEE for giving them the opportunity to organize this special issue. Finally, they would like to thank J. Sun, Senior Publications Editor, and V. Damle, Managing Editor, for their advice, guidance, and support during the special issue preparation process. ■

REFERENCES

[1] J. S. Duncan and N. Ayache, "Medical image analysis: Progress over two decades and the challenges ahead," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 1, pp. 85–106, Jan. 2000.

[2] H. Brody, "Medical imaging," *Nature*, vol. 502, p. S81, Oct. 2013.

[3] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Mar. 2008.

[4] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[5] Alibaba Clouder (blog). (Sep. 2017). *Deep Learning vs. Machine Learning vs. Pattern Recognition*. [Online]. Available: https://www.alibabacloud.com/blog/deep-learning- vs-machine-learning-vs-pattern-recognition_207110

[6] D. Shen, G. Wu, and H. Suk, "Deep learning in medical image analysis," *Annu. Rev. Biomed. Eng.*, vol. 19, pp. 221–248, Jun. 2017.

[7] R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*. Hoboken, NJ, USA: Wiley, 1973.

[8] H. Greenspan, B. van Ginneken, and R. M. Summers, "Deep learning in medical imaging: Overview and future promise of an exciting new technique," *IEEE Trans. Med. Imag.*, vol. 35, no. 5, pp. 1153–1159, May 2016.

[9] G. Litjens *et al.*, "A survey on deep learning in medical image analysis," *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017.

[10] *Space Technology Hall of Fame: Inducted Technologies/1994*, Space Foundation, Colorado Springs, CO, USA, 1994.

[11] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, "Image reconstruction is a new frontier of machine learning," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1289–1296, Jun. 2018.

[12] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *Proc. IEEE*, vol. 98, no. 6, pp. 1045–1057, Jun. 2010.

ABOUT THE GUEST EDITORS

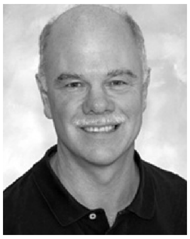

James S. Duncan (Life Fellow, IEEE) received the B.S.E.E. degree (honors) from Lafayette College, Easton, PA, USA, in 1973, the M.S. degree in electrical engineering from the University of California, Los Angeles, CA, USA, in 1975, and the Ph.D. degree in electrical engineering from the University of Southern California, Los Angeles, in 1983.

He has been a Professor of radiology and electrical engineering at Yale University, New Haven, CT, USA, since 1983, where he has also been a Professor of Biomedical Engineering since 2003, the Ebenezer K. Hunt Professor of biomedical engineering since 2007, and a Professor of statistics and data science since 2018. He has served as the Acting Chair and is currently the Director of Undergraduate Studies for biomedical engineering. He has been the Principal Investigator on a number of peer-reviewed grants from the National Institutes of Health and the National Science Foundation over the past 30 years. His research efforts have been in the areas of computer vision, image processing, and medical imaging, with an emphasis on biomedical image analysis. He has published over 270 peer-reviewed articles in these areas.

Dr. Duncan is a Fellow of the American Institute for Medical and Biological Engineering (AIMBE) and the Medical Image Computing and Computer-Assisted Intervention (MICCAI) Society. In 2012, he was elected to the Council of Distinguished Investigators and the Academy of Radiology Research. He received the Enduring Impact Award from the MICCAI Society in 2017. In 2014, he was elected to the Connecticut Academy of Science and Engineering. He has served as the Co-Founder and Co-Editor-in-Chief for *Medical Image Analysis*, as an Associate Editor for the IEEE TRANSACTIONS ON MEDICAL IMAGING, and on the Editorial Board of the *Pattern Analysis and Applications*, the *Journal of Mathematical Imaging and Vision*, and *Modeling in Physiology* of the American Physiological Society. He is a Past President of the MICCAI Society.

Michael F. Insana (Life Fellow, IEEE) received the Ph.D. degree in medical physics from the University of Wisconsin-Madison, Madison, WI, USA, in 1983. His dissertation topic was acoustic scattering and absorption in biological tissues.

He has been Professor of Bioengineering at the University of Illinois at Urbana–Champaign (UIUC), Urbana, IL, USA since 2004. He is currently a Donald Biggar Willett Professor of engineering with joint appointments in Electrical and Computer Engineering and the Beckman Institute for Advanced Science and Technology, UIUC. His research focuses on the development of novel ultrasonic instrumentation and methods for imaging soft tissue microstructure, elasticity, and blood flow. Current projects include developing methods that combine finite-element methods with cooperative neural networks to form images of viscoelastic properties.

Dr. Insana is a Fellow of the American Institute for Medical and Biological Engineering (AIMBE), the Institute of Physics (IoP), and the Acoustical Society of America. He chaired the Gordon Research Conference on Image Science in 2016 and 2018. He was the Editor-in-Chief of the IEEE TRANSACTIONS ON MEDICAL IMAGING from 2015 to 2020.

Nicholas Ayache (Member, IEEE) graduated from the École des Mines, Saint Étienne, France, in 1980, and received the Ph.D. and Habilitation degrees from the University of Paris-Sud, Orsay, France, in 1983 and 1988, respectively.

He was an Invited Scientist with the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA; Harvard University, Cambridge; and the Brigham and Woman's Hospital, Boston, MA, USA, in 2007. He was a Chair Professor with the Collège de France, Paris, France, in 2014, where he introduced a new course on the "Personalized Digital Patient." He served as the Chief Scientific Officer (CSO) with the Institut Hospitalo-Universitaire (IHU), Strasbourg, France, from 2012 to 2015. He is a Research Director of class exceptional at Inria, Sophia Antipolis, France, the Head of the Epione Research Project Team dedicated to e-Patients for e-Medicine, and the Scientific Director of the new artificial intelligence (AI) institute 3IA Côte d'Azur, France. He co-founded seven startup companies in image processing, computer vision, and biomedical imaging. He is a Co-Founder and a Co-Editor-in-Chief of the *Journal of Medical Image Analysis* (Elsevier) (IF $=$ 8.88 in 2019). He is the author and coauthor of more than 400 scientific publications that received 43 000$+$ citations (H-index 103$+$) according to Google Scholar. His current research focuses on AI methods to improve diagnosis, prognosis, and therapy from medical images and clinical, biological, behavioral, and environmental data available on the patient.

Dr. Ayache was a Laureate of the European Research Council (ERC) from 2012 to 2017, the Grand Prize Inria—Académie des sciences in 2014, the Microsoft Prize for Research in Europe (Royal Society and Academy of Sciences) in 2008, and the EADS Foundation Prize in Information Sciences in 2006. He is a member of the French Academy of Sciences and Academy of Surgery. He is a Fellow of the American Institute for Medical and Biological Engineering (AIMBE), the European Alliance of Medical and Biological Engineering and Science (EAMBES), and the Medical Image Computing and Computer Assisted Intervention (MICCAI) Society from which he received the Enduring Impact Award in 2013.