Image Reconstruction Is a New Frontier of Machine Learning

— Editorial for the Special Issue "Machine Learning for Image Reconstruction"

Ge Wang, Jong Chu Ye, Klaus Mueller, Jeffrey A. Fessler

May 2, 2018

1 Introduction

Over past several years, machine learning, or more generally artificial intelligence, has generated overwhelming research interest and attracted unprecedented public attention. As tomographic imaging researchers, we share the excitement from our imaging perspective [1], and organized this special issue dedicated to the theme of "Machine Learning for Image Reconstruction". This special issue is a sister issue of the special issue published in May 2016 of this journal with the theme "Deep Learning in Medical Imaging" [2]. While the previous special issue targeted medical image processing/analysis, this special issue focuses on data-driven tomographic reconstruction. These two special issues are highly complementary, since image reconstruction and image analysis are two of the main pillars for medical imaging. Together we cover the whole workflow of medical imaging: from tomographic raw data/features to reconstructed images and then extracted diagnostic features/readings.

In perspective, computer vision and image analysis are great examples of machine learning, especially deep learning. While computer vision and image analysis deal with existing images and produce features of these images (images to features), tomographic reconstruction produces images of internal structures from measurement data which are various features (line integrals, harmonic components, etc.) of the underlying images (features to images). Recently, machine learning, especially deep learning, techniques are being actively developed worldwide for tomographic reconstruction, as clearly evidenced by the high-quality papers included in this special issue. In addition to well-established *analytic* and *iterative* methods for tomographic image reconstruction, it is now clear that machine learning is an emerging approach for image reconstruction, and image reconstruction is a new frontier of machine learning.

2 Papers Included in the Special Issue

We have 20 papers in this special issue. Each paper was reviewed by 3-4 experts in the area of research and went through a rigorous revision process, composed of typically two rounds of revision. All of the papers in this special issue apply machine learning, especially neural network techniques, for tomographic image reconstruction. There are two major stages in the image reconstruction pipeline: (1) during the reconstruction in the form of forward projection, backprojection, and/or regularization, or (2) after reconstruction as a post-processing operation to mitigate noise and artifacts. Some papers use machine learning to overcome artifacts due to insufficient data, as occurring in a limited angle or sparse data acquisition mode, or image artifacts due to adverse object conditions such as metal in a patient. Also, this special issue covers a number of imaging modalities, such as X-ray CT, MRI, PET, and photoacoustic tomography. In the following, we shall give a narrative of each paper, appropriately grouping them under the multi-pronged categories presented above.

The papers applying data-driven methods in the reconstruction process include the work by Adler and Öktem "*Learned Primal-dual Reconstruction*" [3], which learns the reconstruction operator and combines deep learning with model-based reconstruction. Also, Chen *et al.* in "*LEARN: Learned Experts' Assessment-based Reconstruction Network for Sparse-data CT*" [4] extend sparse coding and learn all regularization terms and parameters in an iteration dependent manner. Würfl *et al.* contribute "*Deep Learning Computed Tomography: Learning Projection-Domain Weights from Image Domain in Limited Angle Problems*" [5], which is a framework for learning the weight and correction matrix and performing cone-beam CT reconstruction. Zheng *et al.* present "*PWLS-ULTRA: An Efficient Clustering and Learning-Based Approach for Low-Dose 3D CT Image Reconstruction*" [6], which uses the penalized weighted least squares (PWLS) method based on an efficient union of learned transforms. Gupta *et al.* in "*CNN-Based Projected Gradient Descent for Consistent CT Image Reconstruction*" [7] replace the projector in a projected gradient descent (PGD) search with a convolutional neural network and apply it also to the sparse-view CT problem. Chen *et al.* in "*Statistical Iterative CBCT Reconstruction Based on Neural Network*" [8] learn the penalty function for statistical iterative reconstruction. Finally, Shen *et al.* present "*Intelligent Parameter Tuning in Optimization-based Iterative CT Reconstruction via Deep Reinforcement Learning*" [9], in which they employ reinforcement learning on-fly to tune parameters for total variation (TV)-regularized CT image reconstruction.

The papers that apply deep learning as an image-space operator are also impressive for the post-reconstruction improvement they were able to achieve. The work by Yang *et al.* "*Low Dose CT Image Denoising Using a Generative Adversarial Network with Wasserstein Distance and Perceptual Loss*" [10] demonstrates a promising combination of traditional reconstruction and network-based denoising for low-dose CT. Independently, Kang *et al.* describe the "*Deep Convolutional Framelet Denoising for Low-Dose CT via Wavelet Residual Network*" [11]. Zhang *et al.* contribute "*A Sparse-View CT Reconstruction Method Based on Combination of DenseNet and Deconvolution*" [12]. Han *et al.* present "*Framing U-Net via Deep Convolutional Framelets: Application to Sparse-view CT*" [13]. Zhang *et al.* address the mitigation of metal artifacts in "*Convolutional Neural Network based Metal Artifact Reduction in X-ray Computed Tomography*" [14] through ensemble learning. Finally, Shan *et al.* investigate "*Low-Dose CT via Transfer Learning from a 2D Trained Network*" [15], using a conveying path based convolutional encoder-decoder (CPCE) network for CT image denoising, in which an initial 3D CPCE denoising model is initialized by extending a trained 2D CNN.

There are three papers that apply deep learning in MRI within a compressed sensing framework. Quan *et al.* present "*Compressed Sensing MRI Reconstruction using a Generative Adversarial Network with a Cyclic Loss*" [16], in which they replace the iterative solver by a faster GAN. Likewise, Yang *et al.* offer "*DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction*" [17]. Gözcü *et al.* describe "*Learning-Based Compressive MRI*" [18], which applies deep learning to optimize MRI sampling patterns. The final set of papers address two more modalities, PET and photoacoustic tomography. Kim *et al.* in "*Penalized PET Reconstruction Using Deep Learning Prior and Local Linear Fitting*" [19] incorporate a deep neural network powered denoising step into an iterative PET reconstruction framework. Yang *et al.* in "*Artificial Neural Network Enhanced Bayesian PET Image Reconstruction*" [20] use a neural network to model a highly nonlinear and spatial-varying patch-wise mapping between a reconstructed image and an enhanced image. Allman *et al.* in "*Photoacoustic Source Detection and Reflection Artifact Removal Enabled by Deep Learning*" [21] employ deep learning techniques to identify reflection artifact stemming from small tips of needles, catheters, and so on for removal in experimental photoacoustic data. And last but not least, Hauptman *et al.* contribute "*Model Based Learning for Accelerated, Limited-View 3D Photoacoustic Tomography*" [22], which uses an iterative reconstruction scheme.

3 Data and Software Used for the Special Issue

A proactive goal of this special issue was to encourage authors to make their codes and data publicly available for reproducible research and resource sharing. Our initiative received very positive responses. By far the most common mechanism for sharing codes has been the development platform GitHub[1]. This site has existed for 10 years and seems likely to persist for many years to come. Whether any given commercial web site will endure as long as a professional organization like IEEE is an open question. One group of authors used the IEEE DataPort repository site[2] for their paper. In terms of software infrastructure, the two most common tools used are TensorFlow with various interfaces[3], and MathWorks MatConvNet[4].

Surveying the papers in this special issue, certain data collections were particularly popular for training, testing and validating data-driven methods. For papers on X-ray CT imaging, the most prevalent data used was that associated with the 2016 AAPM Low Dose CT Grand Challenge[5] [23]. The low-dose sinograms in this collection were prepared using software noise insertion [24]. The importance of access to such data was identified as a key need in the 2011 NIBIB Summit on Low-dose CT [25]. Some CT papers also used the XCAT digital phantom [26, 27] and local datasets. The MRI papers used a wide variety of data sources, including BrainWeb[6] [28] as well as the IXI dataset[7]. These sources seem to be largely MRI images, not k-space data. Apparently, there remains a need for a repository of prospectively fully-sampled and under-sampled multi-coil MRI data.

4 Data-driven Reconstruction as the Third Wave

In retrospect, the field of image reconstruction has experienced three phases/waves of development. The first phase, dating back decades to the inception of each modality, involved *analytical* methods that use an idealized mathematical model of an imaging system. Classical examples are the filtered back-projection method (FBP) for computed tomography (CT) and the inverse Fourier transform for MRI. Typically, these analytical methods consider the sampling properties of the imaging system but few, if any, properties of the object being imaged. These reconstruction methods are extensively used because they are computationally very efficient. The second phase features iterative reconstruction. The iterative methods account for the physical and statistical properties of the imaging system, despite unavoidable mismatches between the mathematical model and numerous physical factors such as non-ideal detector responses and inhomogeneous magnetic fields. Such methods became commercially available for major modalities including SPECT and PET in 1990s, X-ray CT in 2012, and MRI in 2017. Many of these methods are based on statistical object models like Markov random fields or on regularization methods like roughness penalties. In recent years, regularizers based on sparsity and low-rank models using wavelet transforms, finite differences, nuclear norm, manifold learning, and so on have been emphasized. To the best of our knowledge, all commercial image reconstruction methods currently available for PET, CT and MR use regularizers that were designed mathematically, such as the relative difference prior for PET [29, 30], finite difference for CT [31], and combination of wavelets and total variation for MRI [32]. A third phase of image reconstruction research has emerged in the last few years, and involves *data-driven* and *learning-based* methods. These methods

[1] https://github.com
[2] https://ieee-dataport.org
[3] https://www.tensorflow.org
[4] http://www.vlfeat.org/matconvnet
[5] https://www.aapm.org/GrandChallenge/LowDoseCT
[6] http://www.bic.mni.mcgill.ca/brainweb
[7] http://brain-development.org/ixi-dataset

can supplement or replace "human-defined" signal models with counterpart networks learned from big data. Many of the newest methods draw extensively on techniques from the field of machine learning, and it was this wave of algorithmic development that inspired this special issue.

There are various ways to use machine learning techniques for tomographic image reconstruction as discussed in [1]. In some modalities, there are high-quality scans available. Those scans can be used to learn signal models. Then, the learned signal models help reconstruct images from poorer quality data (i.e., under-sampled, photon-limited, or subject to other constraints). A representative method in this class is to learn a dictionary that can represent patches in underlying images with sparse coefficients. Given a set of training images, we can learn the dictionary by solving an optimization problem [33, 34]. An alternative is to learn a sparsifying transform (instead of a dictionary) in the training process [35]. After learning a dictionary or transform from training data, we can reconstruct an unknown image from (low quality and often under-determined) measurements, which is often modeled in the linear form and solved as an optimization problem; see, for example, [34]. A variation of such methods is to learn the dictionary and transform jointly during the reconstruction process; this is called *blind* or *adaptive* learning [33, 34, 36]. Instead of extracting patches, an alternative is to learn convolutional models (filters) from training data [37, 38] in either a synthesis or analysis form. After the filter learning, we can use learned filters to regularize image reconstruction. This is related to convolutional neural networks (CNNs), because it applies a collection of filters to a candidate image.

All of the above cost functions for image reconstruction require iterative methods to perform optimization. One can view such an iterative algorithm as a recurrent neural network involving layers with operations like filtering, thresholding and summing. Instead of treating those operations as fixed components of the iterative algorithm, we can "unroll the loop" of the iterative algorithm into a sequence of operational layers and then optimize the layers in a data-driven fashion. The first method of this type was LISTA (learned ISTA) [39] designed for generic use. Following this seminal work, a number of algorithms were recently developed for image reconstruction, including some papers in this special issue. Some such methods include the system model $\boldsymbol{A}$ in the loop, whereas others attempt to reduce the computational cost by learning filters to approximate the operation $\boldsymbol{A}'\boldsymbol{A}$ better. For medical tomographic imaging, minimizing the cost function typically requires a large amount of computation. One of the potential appeals of some "deep reconstruction" methods is that they may require much less computation after being trained. Typically, such methods start with a crude initial image, and then enhance (denoise, dealias, etc.) a current image using a deep network. Many papers in this issue are in this form.

5 Theoretical Issues Relevant to Deep Reconstruction

Despite the intriguing empirical performance improvement with machine learning techniques over classical analytic and iterative approaches, one of the major hurdles to widespread acceptance of the deep reconstruction or deep imaging approach is that neural networks are still largely black-boxes. The theoretical origins of the success are not well understood yet. In most laboratories worldwide, the specific network architecture and training scheme for each image reconstruction problem is found by trial and error, rather than derived from a governing theory or specific design guidelines. Although a number of interesting theoretical findings and insights [40, 41, 42, 43] have been reported to open the black-box of the neural network, the relevance to image reconstruction applications remains weak, because the theoretical issues in those studies are often remote from specifics of inverse problems and their solutions.

There are two major theoretical issues of machine learning in general, and for image reconstruction in particular: (1) what network architecture should be used for optimal performance (why go deep?), and (2) which neural network training procedure should be applied to guarantee the convergence (how to train?). Most of the existing theoretical results focus on the latter problem, while the design of the network architecture is largely left to experimental exploration.

Recently, progress is being made toward the understanding of the network architecture. For example, the deep convolutional framelets are analyzed in [44]. In this work, the encoder-decoder network emerges from the learning-based Hankel matrix decomposition, and the left and the right bases correspond to the user-defined pooling and trainable convolutional filters. The low-rank Hankel matrix has been successfully used in compressed sensing [45, 46, 47, 48, 49]. This link between the compressed sensing and deep learning problems can help address open network design problems by revealing the role of the filter channels and pooling layers, and can improve the network performance [11, 13]. As another example, the success of deep learning is attributed to not only mathematics but also physics [50]. Although neural networks can approximate most functions in principle, the class of practical functions often stays on a low-dimensional manifold. Fundamental properties in physics such as symmetry, locality, compositionality, and polynomial log-probability translate into exceptionally simple neural networks. When the statistical process generating the data is of a hierarchical form governing by physics, a deep neural network can be more efficient than a shallow one [50]. The "no-flattening theorems" indicate that deep networks cannot be accurately and efficiently approximated by shallow ones [50]. Moreover, although it is alleged that the nonlinearity of the rectified linear unit (ReLU) allows the conical decomposition of the convolution framelet basis by enforcing the positivity of the frame coefficients, the high dimensional geometric understanding of the multilayer conic decomposition is not fully understood. In addition, the arguments behind the need to use skipped compounds must be reviewed with more empirical evidence and mathematical analysis.

Of particular practical relevance, for a given network architecture the issue of convergence to the global minimizer has been an extensively studied theoretical topic [40, 41, 42, 43]. The recent theoretical advances in non-convex optimization (see [51] and references therein) bring valuable tools, and the so-called optimization landscape [52, 53] is a key for such studies. For example, the authors in [41, 42, 43] showed that a well-designed neural network has no spurious local minimizers so that the gradient descent optimizer converges to the global minimizer. In addition, the theory of information bottleneck [40] provides an important dynamic view to explain the solution trajectory during the training process. However, the existing theoretical results for the convergence are mainly focused on simple network architectures for classification. Therefore, it will be an important research direction for the imaging community to investigate how this optimization landscape changes with different neural network architectures under tomographic data constraints and domain-specific priors.

6 Research Opportunities in Data-driven Imaging

In addition to the above-mentioned theoretical research, machine learning methods are constantly introduced to solve inverse problems, as reported in this special issue and many other papers. The evolution has been rapid with unprecedented challenges and opportunities, in the following six distinct but inter-related categories: (1) big data generation, (2) image domain learning [13, 54, 55, 56, 57, 58, 59, 60, 61], (3) data domain learning [5, 62, 63], (4) hybrid reconstruction schemes [3, 4, 7, 11, 16, 54, 64, 65, 66, 67, 68, 69, 70], (5) end-to-end workflows [71, 72], and (6) convergence of deep imaging and other emerging technologies.

First, machine learning based image reconstruction demands a large amount of curated, calibrated, and preferably annotated data for training, validation, tuning, and testing. In reference to machine learning successes with computer vision or image processing, big data such as *ImageNet* [73] has millions of images. Due to the complexity of the tomographic mapping between sinogram and image spaces, the big data size for CT as an example is likely comparable to or larger than that of *ImageNet*. Then, a major challenge is lack of data for several reasons. Despite a huge amount of data in existence worldwide, only a tiny fraction is available for research due to privacy, legal, and business-related concerns. As an example, CT raw data are tightly protected by companies and generally inaccessible to researchers. Also, a research project normally has limited resources and only targets a specific anatomical site/disease, and it is difficult to collect sufficiently big data as compared to the well-known benchmark such as *ImageNet* [73]. Furthermore, at the developmental stage, a large amount of annotated data are not in existence at all, such as in the cases of future cardiac CT with high-resolution and photon-counting detectors. As a consequence, data sizes of hundreds or even less are typically seen in manuscripts submitted to this special issue. A limited data source may cause overfitting, inaccuracy, artifacts, etc. [2], impeding the advancement and translation of data-driven tomographic reconstruction research. Perhaps that sufficiently representative tomographic data can be obtained by intelligent augmentation with coupling of simulation data, emulation data and transfer learning.

Second, low-hanging fruits for machine learning oriented medical imaging are to design a deep network in the image domain. Noisy or artifact-corrupted images are generated from measurement data. A neural network can be trained to learn the artifacts from big data. For example, the low-dose and sparse CT neural networks [13, 56, 57, 58, 60, 61] are often in this style. Similarly, the earlier applications of the neural networks for compressed MRI were designed to remove aliasing artifacts after obtaining a Fourier inversion image from down-sampling k-space data [59, 74]. A key benefit of these image domain algorithms is that off-the-shelf tools from the computer vision literature can be adapted to enhance image quality. Interestingly, this approach also leads to skepticism as to whether an improvement in image quality is actual or just a cosmetic change. In [1], it was suggested that "with deep neural networks, depth and width can be combined to efficiently represent functions with a high precision, and perform powerful multi-scale analysis, quite like wavelet analysis but in a nonlinear manner." Mathematically, the authors of [44] show that the architecture of the neural network is actually a signal representation similar to the representation using wavelets or framelets [75]. The main difference is that the bases are learned from the training data. Hence, the image-based neural network can be understood as an image noise removal algorithm, even better than the wavelet or framelet shrinkage widely used in the medical imaging community. This suggests that the image enhancement from the neural network is based on the well-established signal processing principle, and the resultant improvement is intrinsic and not just cosmetic.

Third, a more radical strategy is to make the full use of the measurement data with a deep neural network [1]. A good example is the so-called AUTOMAP (which is an automated transform through manifold approximation) for MRI [62]. Featured in a high profile journal, AUTOMAP directly learns the network-based mapping between the measurement and the image. The key difference of this architecture from other CNN setups is that the first layer is fully-connected, since k-space data are coefficients of the Fourier transform which is a global transform. Then, MRI image reconstruction from under-sampled k-space data can be solved by coupling a fully-connected layer with other layers all of which are trained to invert the Fourier transform. While the idea of learning the inverse mapping using a fully-connected layer is intriguing, a major drawback of AUTOMAP is its huge memory requirement to store the fully-connected layer. However, for relatively small size images, AUTOMAP gives a promising direction of research that directly related the measurement data to images through neural network. Another potential criticism is that the neural network may not need to learn the Fourier transformation from scratch, since a memory-efficient analytic transform is already available. The recent efforts in training the sinogram domain filter using a neural network offers an important clue toward that direction [5]. In this paper, similar to AUTOMAP, the neural network was realized in the measurement domain, and trained to map from the sinogram to the image. The training goal is to estimates the data-driven ramp-type filter by minimizing the image domain loss. However, in contrast to AUTOMAP, a fully-connected layer is not necessary. While the performance improvement of this approach over the image domain neural network needs further verification, the full utilization of measurement domain data with neither iterations nor fully-connected layers is attractive and deserves further investigation. Another angle to do data-based learning is to improve tomographic raw data themselves. In [76], Feng *et al.* develop a machine learning method to correct for the pileup effect of photon-counting CT data. This model-free and data-driven method is a fully connected cascade artificial neural network trainable with measurements and true counts for high fidelity energy-sensitive data, which should be a better starting point for data-domain learning.

Fourth, hybrid reconstruction methods integrate merits of data- and image-domain learning methods. One of the earliest studies is the variational neural network designed by Hammernik *et al.* [54]. In that context, the image prior is modeled as a combination of unknown nonlinearites and weights that can be estimated with training data. To find these parameters, the gradient update is performed in multiple steps that map directly to the layers of a neural network. Similar methods were used in ADMM-Net [64], in which the unfolded steps of the alternating directional method of multiplier (ADMM) algorithm are mapped into layers of a neural network. For dynamic cardiac MRI, Schlemper *et al.* [66] proposed a cascaded CNN, in which the network alternates repetitions between the data consistency layers and the image domain denoising layers. Quan *et al.* [16] employed the cyclic consistency in the k-space and the image space to solve a compressed sensing MRI problem. Another way is to use neural network as a prior model within a model-based iterative reconstruction (MBIR) framework. The earliest form of this idea was proposed by Wang *et al.* [65], in which the CNN-based prior was put in a formulation of compressed sensing MRI. Also, using the plug-and-play approach the denoising step of the ADMM was replaced with a neural network denoiser [69, 70]. A more sophisticated form of such an approach is to make a neural network as a projector for a desirable functional space [7]. Similarly, based on the observation that a neural network can be interpreted as a framelet signal representation with its shrinkage behavior controlled by filter channels [44], a framelet-based denoising algorithm was suggested that iterates between the data consistency step and the neural network shrinkage step [11]. The main advantage of these algorithms is their provable convergence, since the algorithms depend heavily on the proximal optimization. A recent study by Adler *et al.* [3] established an elegant framework using trainable primal and dual steps that guarantees the convergence. The idea can be further extended using the recursive neural network (RNN), in which relatively shallow neural networks are used as a prior model for parallel MRI [68]. Thanks to the full use of the measurement data, these algorithms offer consistent improvement over the image domain counterparts. However, one of the drawbacks is that computational advantages of the feed-forward neural network are lost in these iterative frameworks. For example, in CT and PET reconstructions with neural network priors [69, 70], multiple projections and backprojections become necessary, taking substantial computational time. However, 3D iterative methods are already available commercially for CT and PET for many years now so reducing computing time may be less urgent than improving image quality.

Fifth, the end-to-end workflow has a significant potential that will integrate deep reconstruction and radiomics for optimal diagnostic performance [71, 72]. Radiomics, a hot area for years, utilizes extensive features mined from images using sophisticated algorithms including deep neural networks. A synergistic opportunity exists between deep reconstruction and radiomics, i.e., the unification of tomographic reconstruction and radiomics for what we call "*rawdiomics*" (*raw d*ata/*i*nfo+*omics*) so that the space of features can be widened for better diagnostic performance. Driven by the desire to train the deep neural network directly with tomographic data, Wu *et al.* proposed an end-to-end network for lung CT nodule detection [72]. Their network is the unrolled version of an iterative reconstruction process from low-dose CT data through images to final diagnosis. The involved reconstruction and analysis parts of the overall network were jointly trained, yielding better sensitivity and accuracy than what can be obtained when reconstruction and analysis were separately done [72].

Lastly and a bit speculative, given unprecedented progresses in the engineering field over the past decade or so, cutting-edge medical imaging, machine learning, robot, high-performance computing, internet, and auto-driving technologies could be combined to change the landscape of the medical imaging world. We envision a paradigm shift in medical imaging, from hospital/clinic/center-oriented to mobile, intelligent, integrated, and cost-effective services promptly delivered wherever and whenever needed. This patient-oriented imaging service will be not only convenient (an analogy is telephone booths versus smart phones) but also cost-effective (highly automated process consisting of an auto-driving scanner and a robotic technician, who can come to your place as a Uber taxi can) and even necessary in natural disaster scenes and near battle fields, which should be a good example of what we call *"Internet of Imaging Services"*.

7 Tips for Interested Novice Researchers

In this section we aim to offer some advice for those peers/students in our field who are not involved in machine learning based medical imaging yet but have read this editorial due to their interest in the emerging area.

Machine learning, generally speaking, is a mechanism that bypasses painstaking analytical model construction via an equation by instead learning these models from data. It identifies patterns that can be arbitrarily complex and enables knowledge-finding and decision-making without humans in the loop. While machine learning in principle is not new, there are two enabling developments that have helped it to flourish, namely, (1) the widespread availability of big data which can be stored on cheap disks and communicated over high-bandwidth networks; and (2) the tremendous advances in computer hardware, algorithms and software which can run the compute-intensive machine learning algorithms within practically tolerable time and generate results outperforming human experts in many cases. Machine learning can be supervised, semi-supervised, or unsupervised. Examples for supervised algorithms are learning a classifier or a regression model, while examples for unsupervised machine learning are clustering or the derivation of mixture models.

Deep Learning is the currently most active branch of machine learning inspired by the structure and function of the brain, simulated in artificial neural networks. Neural networks with backpropagation learning have been around for many years, but the realization that one can build deep networks that go from low-level feature recognition to abstract concepts, layer by layer, has revolutionized the field [77]. Also, deep learning has crucially benefited from the development of high performance commodity parallel hardware such as the GPU (Graphics Processing Unit) and more recently also custom hardware developed by major industrial players in the field, such as Google's TPU (Tensor Processing Units). Deep neural network architects usually make use of highly versatile open source packages such as TensorFlow[8], Theano[9], Caffe[10], Keras[11], PyTorch[12], and various others[13].

Readers who seek to come up to speed with the topics of machine learning may begin with the Wikipedia page of the same name[14] and follow some of links to read on further. Then, one may consider taking one of the quite numerous and diverse Coursera courses[15]. The one by Stanford University taught by Andrew Ng was the first, and in fact one of the first Coursera courses overall.

[8] https://www.tensorflow.org/
[9] http://deeplearning.net/software/theano
[10] http://caffe.berkeleyvision.org/
[11] https://keras.io/
[12] https://pytorch.org/
[13] https://en.wikipedia.org/wiki/Comparison_of_deep_learning_software
[14] https://en.wikipedia.org/wiki/Machine_learning
[15] https://www.coursera.org/browse/data-science/machine-learning?languages=en

For deep learning, there are quite a few online tutorials. The website "Neural networks and deep learning" has a wealth of material and at the same time is quite intuitive to read[16]. The MIT Press book [78] by Goodfellow, Bengio, and Courville is online and an excellent read[17]. Moreover, the source codes and data set provided by the authors in this special issue are believed to be good starting points to understand how deep learning algorithms can be implemented for image reconstruction purpose. Finally, the folks from Google Brain and other groups put together a highly interactive website that focuses on the interpretability of deep neural networks[18].

8 Conclusion

We feel humbled and privileged to be given the opportunity of editing this special issue. We are obligated to express our appreciation for the excellent jobs done by numerous reviewers, outstanding infrastructural support by the TMI office staff especially Ms. Deborah Insana, important guidance by the Editor-in-Chief Dr. Michael Insana, and the approval of our initiative by the TMI Steering Committee for this special issue. Needless to say, we have learnt a great deal from all the submissions by the authors who are among the most proactive and creative colleagues in our community. Without any of these, it would have been impossible to finish this special issue which we hope to have a major and lasting value.

In conclusion, big data, machine learning and artificial intelligence will fundamentally impact the field of medical imaging. An aggressive view is that "machine learning will transform radiology significantly within the next 5 years" [79]. Regardless the pace at which machine learning is being translated to hospitals and clinics, the future for imaging research, development and applications seems definitely exciting for us and younger generations.

References

[1] G. Wang, "A perspective on deep imaging," *IEEE Access*, vol. 4, pp. 8914–8924, 2016.

[2] H. Greenspan, B. van Ginneken, and R. M. Summers, "Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique," *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1153–1159, 2016.

[3] J. Adler and O. Öktem, "Learned primal-dual reconstruction," *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[4] H. Chen, Y. Zhang, Y. Chen, J. Zhang, W. Zhang, H. Sun, Y. Lv, P. Liao, J. Zhou, and G. Wang, "LEARN: Learned experts' assessment-based reconstruction network for sparse-data CT," *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[5] T. Würfl, M. Hoffmann, V. Christlein, K. Breininger, Y. Huang, M. Unberath, and A. K. Maier, "Deep learning computed tomography: Learning projection-domain weights from image domain in limited angle problems," *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[6] X. Zheng, S. Ravishankar, Y. Long, and J. A. Fessler, "PWLS-ULTRA: An efficient clustering and learning-based approach for low-dose 3D CT image reconstruction," *IEEE Transactions on Medical Imaging (this issue)*, 2017.

[16] http://neuralnetworksanddeeplearning.com/
[17] http://www.deeplearningbook.org
[18] https://distill.pub/2018/building-blocks/

[7] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser, “CNN-based projected gradient descent for consistent image reconstruction,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[8] B. Chen, K. Xiang, Z. Gong, J. Wang, and S. Tan, “Statistical iterative CBCT reconstruction based on neural network,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[9] C. Shen, Y. Gonzalez, L. Chen, S. B. Jiang, and X. Jia, “Intelligent parameter tuning in optimization-based iterative CT reconstruction via deep reinforcement learning,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[10] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun, and G. Wang, “Low dose CT image denoising using a generative adversarial network with Wasserstein distance and perceptual loss,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[11] E. Kang, W. Chang, J. Yoo, and J. C. Ye, “Deep convolutional framelet denosing for low-dose CT via wavelet residual network,” *IEEE Transactions on Medical Imaging(this issue)*, 2018.

[12] Z. Zhang, X. Liang, X. Dong, Y. Xie, and G. Cao, “A sparse-view CT reconstruction method based on combination of DenseNet and deconvolution,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[13] Y. Han and J. C. Ye, “Framing U-Net via deep convolutional framelets: Application to sparse-view CT,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[14] Y. Zhang and H. Yu, “Convolutional neural network based metal artifact reduction in x-ray computed tomography,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[15] H. Shan, Y. Zhang, Q. Yang, U. Kruger, M. K. Kalra, L. Sun, W. Cong, and G. Wang, “3D convolutional encoder-decoder network for low-dose CT via transfer learning from a 2D trained network,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[16] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, “Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[17] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, *et al.*, “Dagan: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[18] B. Gözcü, R. K. Mahabadi, Y.-H. Li, E. Ilıcak, T. Cukur, J. Scarlett, and V. Cevher, “Learning-based compressive MR,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[19] K. Kim, D. Wu, K. Gong, J. Dutta, J. H. Kim, Y. D. Son, H. K. Kim, G. E. Fakhri, and Q. Li, “Penalized PET reconstruction using deep learning prior and local linear fitting,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[20] B. Yang, L. Ying, and J. Tang, “Artificial neural network enhanced bayesian PET image reconstruction,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[21] D. Allman, A. Reiter, and M. A. L. Bell, “Photoacoustic source detection and reflection artifact removal enabled by deep learning,” *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[22] A. Hauptmann, F. Lucka, M. Betcke, N. Huynh, J. Adler, B. Cox, P. Beard, S. Ourselin, and S. Arridge, "Model based learning for accelerated, limited-view 3D photoacoustic tomography," *IEEE Transactions on Medical Imaging (this issue)*, 2018.

[23] C. H. McCollough, A. C. Bartley, R. E. Carter, B. Chen, T. A. Drees, P. Edwards, D. R. Holmes, A. E. Huang, F. Khan, S. Leng, K. L. McMillan, G. J. Michalak, K. M. Nunez, L. Yu, and J. G. Fletcher, "Low-dose CT for the detection and classification of metastatic liver lesions: Results of the 2016 Low Dose CT Grand Challenge," *Med. Phys.*, vol. 44, pp. e339–52, Oct. 2017.

[24] L. Yu, M. Shiung, D. Jondal, and C. H. McCollough, "Development and validation of a practical lower-dose-simulation tool for optimizing computed tomography scan protocols," *J. Comp. Assisted Tomo.*, vol. 36, pp. 477–87, July 2012.

[25] C. H. McCollough, G. H. Chen, W. Kalender, S. Leng, E. Samei, K. Taguchi, G. Wang, L. Yu, and R. I. Pettigrew, "Achieving routine submillisievert CT scanning: report from the summit on management of radiation dose in CT," *Radiology*, vol. 264, pp. 567–80, Aug. 2012.

[26] W. P. Segars, M. Mahesh, T. J. Beck, E. C. Frey, and B. M. W. Tsui, "Realistic CT simulation using the 4D XCAT phantom," *Med. Phys.*, vol. 35, pp. 3800–8, Aug. 2008.

[27] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, "4D XCAT phantom for multimodality imaging research," *Med. Phys.*, vol. 37, pp. 4902–15, Sept. 2010.

[28] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans, "Design and construction of a realistic digital brain phantom," *IEEE Trans. Med. Imag.*, vol. 17, pp. 463–8, June 1998.

[29] S. Ahn, S. G. Ross, E. Asma, J. Miao, X. Jin, L. Cheng, S. D. Wollenweber, and R. M. Manjeshwar, "Quantitative comparison of OSEM and penalized likelihood image reconstruction using relative difference penalties for clinical PET," *Phys. Med. Biol.*, vol. 60, pp. 5733–52, Aug. 2015.

[30] J. Nuyts, D. Beque, P. Dupont, and L. Mortelmans, "A concave prior penalizing relative differences for maximum-a-posteriori reconstruction in emission tomography," *IEEE Trans. Nuc. Sci.*, vol. 49, pp. 56–60, Feb. 2002.

[31] J.-B. Thibault, K. Sauer, C. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," *Med. Phys.*, vol. 34, pp. 4526–44, Nov. 2007.

[32] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Mag. Res. Med.*, vol. 58, pp. 1182–95, Dec. 2007.

[33] S. Ravishankar and Y. Bresler, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. Med. Imag.*, vol. 30, pp. 1028–41, May 2011.

[34] Q. Xu, H. Yu, X. Mou, L. Zhang, J. Hsieh, and G. Wang, "Low-dose X-ray CT reconstruction via dictionary learning," *IEEE Trans. Med. Imag.*, vol. 31, pp. 1682–97, 9 2012.

[35] S. Ravishankar and Y. Bresler, "Learning sparsifying transforms," *IEEE Trans. Sig. Proc.*, vol. 61, pp. 1072–86, Mar. 2013.

[36] S. Ravishankar and Y. Bresler, "Efficient blind compressed sensing using sparsifying transforms with convergence guarantees and application to MRI," *SIAM J. Imaging Sci.*, vol. 8, no. 4, pp. 2519–57, 2015.

[37] B. Wohlberg, "Efficient algorithms for convolutional sparse representations," *IEEE Trans. Im. Proc.*, vol. 25, pp. 301–15, Jan. 2016.

[38] I. Y. Chun and J. A. Fessler, "Convolutional dictionary learning: acceleration and convergence," *IEEE Trans. Im. Proc.*, vol. 27, pp. 1697–712, Apr. 2018.

[39] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proc. Intl. Conf. Mach. Learn*, 2010.

[40] N. Tishby and N. Zaslavsky, "Deep learning and the information bottleneck principle," in *Information Theory Workshop (ITW), 2015 IEEE*, pp. 1–5, IEEE, 2015.

[41] R. Ge, J. D. Lee, and T. Ma, "Learning one-hidden-layer neural networks with landscape design," *arXiv preprint arXiv:1711.00501*, 2017.

[42] S. S. Du and J. D. Lee, "On the power of over-parametrization in neural networks with quadratic activation," *arXiv preprint arXiv:1803.01206*, 2018.

[43] S. S. Du, J. D. Lee, Y. Tian, B. Póczos, and A. Singh, "Gradient descent learns one-hidden-layer CNN: don't be afraid of spurious local minima," *arXiv preprint arXiv:1712.00779*, 2017.

[44] J. C. Ye, Y. Han, and E. Cha, "Deep convolutional framelets: A general deep learning framework for inverse problems," *SIAM Journal on Imaging Sciences*, vol. 11, no. 2, pp. 991–1048, 2018.

[45] K. H. Jin, D. Lee, and J. C. Ye, "A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix," *IEEE Transactions on Computational Imaging*, vol. 2, no. 4, pp. 480–495, 2016.

[46] K. H. Jin and J. C. Ye, "Annihilating filter-based low-rank Hankel matrix approach for image inpainting," *IEEE Transactions on Image Processing*, vol. 24, no. 11, pp. 3498–3511, 2015.

[47] D. Lee, K. H. Jin, E. Y. Kim, S.-H. Park, and J. C. Ye, "Acceleration of MR parameter mapping using annihilating filter-based low rank Hankel matrix (aloha)," *Magnetic resonance in medicine*, vol. 76, no. 6, pp. 1848–1864, 2016.

[48] J. Lee, K. H. Jin, and J. C. Ye, "Reference-free single-pass EPI Nyquist ghost correction using annihilating filter-based low rank Hankel matrix (ALOHA)," *Magnetic resonance in medicine*, vol. 76, no. 6, pp. 1775–1789, 2016.

[49] K. H. Jin, J.-Y. Um, D. Lee, J. Lee, S.-H. Park, and J. C. Ye, "MRI artifact correction using sparse+ low-rank decomposition of annihilating filter-based Hankel matrix," *Magnetic resonance in medicine*, vol. 78, no. 1, pp. 327–340, 2017.

[50] H. W. Lin, M. Tegmark, and D. Rolnick, "Why does deep and cheap learning work so well?," *Journal of Statistical Physics*, vol. 168, no. 6, pp. 1223–1247, 2017.

[51] Y. Chen, Y. Chi, J. Fan, and C. Ma, "Gradient descent with random initialization: Fast global convergence for nonconvex phase retrieval," *arXiv preprint arXiv:1803.07726*, 2018.

[52] J. Sun, Q. Qu, and J. Wright, "A geometric analysis of phase retrieval," in *Information Theory (ISIT), 2016 IEEE International Symposium on*, pp. 2379–2383, IEEE, 2016.

[53] J. Sun, Q. Qu, and J. Wright, "When are nonconvex problems not scary?," *arXiv preprint arXiv:1510.06096*, 2015.

[54] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll, "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic resonance in medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.

[55] K. Kwon, D. Kim, and H. Park, "A parallel MR imaging method using multilayer perceptron," *Medical physics*, vol. 44, no. 12, pp. 6209–6224, 2017.

[56] E. Kang, J. Min, and J. C. Ye, "A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction," *Medical physics*, vol. 44, no. 10, 2017.

[57] H. Chen, Y. Zhang, W. Zhang, P. Liao, K. Li, J. Zhou, and G. Wang, "Low-dose CT via convolutional neural network," *Biomedical Optics Express*, vol. 8, no. 2, pp. 679–694, 2017.

[58] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.

[59] Y. S. Han, J. Yoo, and J. C. Ye, "Deep learning with domain adaptation for accelerated projection reconstruction MR," *Magnetic Resonance in Medicine, https://doi.org/10.1002/mrm.27106*, 2018.

[60] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang, "Low-dose CT with a residual encoder-decoder convolutional neural network," *IEEE transactions on medical imaging*, vol. 36, no. 12, pp. 2524–2535, 2017.

[61] J. M. Wolterink, T. Leiner, M. A. Viergever, and I. Išgum, "Generative adversarial networks for noise reduction in low-dose CT," *IEEE transactions on medical imaging*, vol. 36, no. 12, pp. 2536–2545, 2017.

[62] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, p. 487, 2018.

[63] Y. H. Yoon, S. Khan, J. Huh, and J. C. Ye, "Deep learning in RF sub-sampled B-mode ultrasound imaging," *arXiv preprint arXiv:1712.06096*, 2017.

[64] J. Sun, H. Li, Z. Xu, *et al.*, "Deep ADMM-Net for compressive sensing MRI," in *Advances in Neural Information Processing Systems*, pp. 10–18, 2016.

[65] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, and D. Liang, "Accelerating magnetic resonance imaging via deep learning," in *Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on*, pp. 514–517, IEEE, 2016.

[66] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE transactions on Medical Imaging*, vol. 37, no. 2, pp. 491–503, 2018.

[67] D. Lee, J. Yoo, S. Tak, and J. Ye, "Deep residual learning for accelerated MRI using magnitude and phase networks," *IEEE Transactions on Biomedical Engineering*, 2018.

[68] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model based deep learning architecture for inverse problems," *arXiv preprint arXiv:1712.02862*, 2017.

[69] D. Wu, K. Kim, G. El Fakhri, and Q. Li, "Iterative low-dose CT reconstruction with priors trained by artificial neural network," *IEEE transactions on medical imaging*, vol. 36, no. 12, pp. 2479–2486, 2017.

[70] K. Gong, J. Guan, K. Kim, X. Zhang, G. E. Fakhri, J. Qi, and Q. Li, "Iterative PET image reconstruction using convolutional neural network representation," *arXiv preprint arXiv:1710.03344*, 2017.

[71] M. Kalra, G. Wang, and C. G. Orton, "Radiomics in lung cancer: Its time is here," *Medical physics*, vol. 45, no. 3, pp. 997–1000, 2017.

[72] D. Wu, K. Kim, B. Dong, and Q. Li, "End-to-end abnormality detection in medical imaging," *arXiv preprint arXiv:1711.02074*, 2017.

[73] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255, IEEE, 2009.

[74] D. Lee, J. Yoo, and J. C. Ye, "Deep residual learning for compressed sensing MRI," in *Biomedical Imaging (ISBI 2017), 2017 IEEE 14th International Symposium on*, pp. 15–18, IEEE, 2017.

[75] I. Daubechies, B. Han, A. Ron, and Z. Shen, "Framelets: MRA-based constructions of wavelet frames," *Applied and computational harmonic analysis*, vol. 14, no. 1, pp. 1–46, 2003.

[76] R. Feng, D. Rundle, and G. Wang, "Neural-networks-based photon-counting data correction: Pulse pileup effect," *arXiv preprint arXiv:1804.10980*, 2018.

[77] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.

[78] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*, vol. 1. MIT press Cambridge, 2016.

[79] G. Wang, M. Kalra, and C. G. Orton, "Machine learning will transform radiology significantly within the next 5 years," *Medical physics*, vol. 44, no. 6, pp. 2041–2044, 2017.