Artificial Intelligence for PET Image Reconstruction

Andrew J. Reader¹ and Georg Schramm²

¹*School of Biomedical Engineering and Imaging Sciences, King's College, London, United Kingdom; and* ²*Division of Nuclear Medicine, Department of Imaging and Pathology, KU/UZ, Leuven, Belgium*

Artificial intelligence (AI) continues to deliver a remarkable impact on numerous and highly diverse fields, such as physics, natural language processing, finance, human resources, image processing, protein folding (1), and prediction of viral mutations (2). In broad terms, AI is any technology that can learn how to perform tasks from example data or experiences. This technology contrasts with the conventional paradigm of a human programmer or engineer providing extensive and exhaustive instructions in order for a task to be performed.

The power of AI is beyond question, but its adoption, as with other groundbreaking technologies, can initially lead to concerns, skepticism, and even ethical questions. In particular, use of AI in medical imaging has demonstrated immense potential (3), but a key question is how much we can trust AI in the formation of images that inform clinical decisions, when lives of patients are often at stake.

This brief article will consider the methodologies, benefits, and concerns regarding AI for the case of the formation, or reconstruction, of PET images (4) and will focus on a subdiscipline of AI, namely deep learning (5). We will define deep learning and then use this term interchangeably with AI.

UNDERSTANDING AI AND DEEP LEARNING

So what is deep learning exactly? Deep learning can be considered as a sequence of steps that operate on input data to perform a desired task, with the steps being learned from example inputs and desired outputs (training data). These sequences of operations are comparable to conventional computer code, which similarly executes a sequence of operations designed (without training data) to specifically accomplish tasks. Therefore, deep learning can be more generally regarded as a data-informed, trainable version of our existing, well-established algorithms.

Taking the example task of PET image reconstruction, algorithms that have been developed by the PET reconstruction community over many decades (drawing on knowledge from imaging physics, mathematics and statistics), can now also be integrated into the learning AI paradigm. Better still, state-of-the-art image reconstruction methods can likely be made even more reliable with AI-informed refinement.

However, AI has been frequently misunderstood, either because of the notion of AI being a black box, or as a result of conventional low-dimensional mathematical perspectives on fitting models to

limited data. The black box misconception originates partly from the highly successful use of deep learning in computer vision tasks, in which its performance has launched deep learning to its deserved level of current recognition. Early successes via the automated hierarchical feature-learning of convolutional neural networks have resulted in large uptake of these networks to other tasks, in which there has been a temptation to use these large architectures without careful design considerations, relying instead on large numbers of trainable parameters. Use of poorly justified and highly parameterized architectures has made it easy to dismiss any chance of understanding (let alone designing) these sophisticated nonlinear mappings, fueling AI skepticism. As for conventional mathematical perspectives on the feasibility of optimization and fitting to limited data in high dimensions, these have proven not to be the showstoppers that they were expected to be. On the contrary, deep learning's success has revealed a need to revise our thinking on optimization, regularization, and generalization.

Hence, the rapid progress of AI methods, sometimes with loss of principled design choices and often to the surprise of conventional mathematical thinking, has resulted in concern over the interpretability and trustworthiness of AI. This situation has not been helped by reduced levels of rigor arising from the surge of innovation and exciting successes. But black box concerns (Fig. 1) and conventional mathematical views on optimization are becoming dated perspectives, particularly in the context of deep learning for signal and image processing. In these fields, increasingly meaningful design choices are being made by embedding the AI paradigm into conventional and well-understood algorithmic processing (such as the discrete Fourier and Radon transforms).

WHY USE AI FOR PET IMAGE RECONSTRUCTION?

In applying AI to image reconstruction for PET, we are recognizing that PET image reconstruction actually needs help. First, improving spatial resolution and lowering noise in PET images will very likely assist in the clinical utility of PET. Second, even if current image quality is deemed acceptable, the desire for shorter acquisition times or reduced radiation doses will require more advanced techniques to try and retain standard image quality from lower-count (noisier) data. Similarly, achieving higher temporal resolution, such as for improved motion correction, will likewise demand improved reconstruction.

Let us now recall what reconstruction actually is: it is the use of raw list-mode or projection data acquired from a PET scan to form an image representing a radiotracer's spatiotemporal distribution within the human body. For conventional PET, the spatial resolution of such images is of the order of a few millimeters, and the temporal

Received May 31, 2021; revision accepted Jun. 14, 2021.

For correspondence or reprints, contact Georg Schramm (georg.schramm@kuleuven.be).

Published online July 8, 2021

COPYRIGHT © 2021 by the Society of Nuclear Medicine and Molecular Imaging.

FIGURE 1. AI in PET reconstruction as seen from various perspectives (15).

resolution is of the order of many seconds. These limitations are due to limited photon counts, scanner design, and physics. Nonetheless, advances in statistical image reconstruction methods for PET have made greater use of the acquired data, lowering image noise and improving spatial and temporal resolution, through accurate modeling of the imaging physics and statistics and through use of prior information (including from CT or MRI). Even with such progress, the limited counts and resolution still place a performance ceiling on the potential of PET for clinical imaging, and as mentioned, the desire to reduce the dose and to shorten scan times means that limited data pose ongoing challenges to PET image reconstruction.

This is where AI can make a huge difference, in 2 main ways. First, with sufficient example data, AI can learn the vast (but nonetheless highly restricted) set of PET images that can realistically ever be expected from a PET scan (this set is often referred to as the manifold). For example, we know a PET scan can never deliver a CT or MR image, let alone a natural photographic image. Yet the mathematics of current image reconstruction methods do not exploit any of this obviously robust prior information but instead can readily accommodate wrong images. This is because current state-of-the-art image reconstruction uses simple, mathematically convenient priors for PET images, which are excessively general (e.g., requiring only that the images be smooth, to suppress noise but at the cost of resolution and details). This process discards considerable amounts of a priori information. In contrast, AI's learning of the manifold of all feasible PET images can be applied to make better use of each and every acquired count in a PET scan. Acquired PET data can therefore be projected, or encoded, into this realistic manifold.

Second, since this learned manifold of all feasible PET images can in fact be represented in infinitely many ways, AI can learn how to encode the acquired PET scan data into latent feature representations that best serve our desired goals. These representations include reduced-dimension representations (bottlenecks) to assist in noise reduction and can also involve projection to higher dimensions to assist in classification tasks. The point is that AI can learn how best to capture and encode key explanatory information, salient to our task, from a given scan.

Therefore, the power of AI is not only its ability to learn how to encode into useful latent representations or feature maps, and learn transforms between them, but also its ability to learn how to decode from these latent representations, to generate outputs for various desired tasks. This could be generation of low-noise reconstructed PET images with high resolution, generation of radiological reports, or indeed diagnostic and prognostic predictions. Learning encodings of acquired PET scan data into contextually rich feature spaces consistent with the PET manifold, and decoding into task-specific forms, is the sublimely powerful ability of AI, which PET would do well to exploit more fully.

HOW CAN WE USE AI IN PET IMAGE RECONSTRUCTION?

There are currently 3 main approaches to using AI in PET reconstruction. The first group of approaches, direct AI (e.g., AUTOMAP (6) or DeepPET (7)), learns an encoding from the raw data, via a latent feature space, to decode to the desired image. The key point here is that the overall mapping is trained by supervised learning, in order to take noisy raw PET data and deliver inferences of the ground-truth object or high-quality reference image, according to the pairings of datasets used in the training phase. Direct AI can easily be understood by comparison to conventional curve-fitting and regression tasks, except that in the case of deep learning of PET reconstruction we are performing regressions with extremely high-dimensional vectors. The input raw PET data are fully 3-dimensional sets of measured (time-of-flight) sinograms (with $\sim 10^8$ – 10^9 bins), for mapping to output 3-dimensional images (with $\sim 10^7$ voxels). At present, these direct deep learning methods look to be impractical, having been demonstrated only for small 2-dimensional reconstructions (e.g., 128×128 images), as they have colossal demands for computational memory and training set sizes ($>10^5$ datasets). Furthermore, they may not generalize well for unseen data (e.g., for data that are too different from the example training data). Early tests of direct methods for real-data 2-dimensional PET reconstructions have delivered images that have yet to convince some experts.

By far the more promising methods, sometimes called physics-informed AI, take the learning paradigm from AI and integrate this into our existing state-of-the-art statistical iterative image reconstruction methods. Here, the standard iterative loop of an image reconstruction algorithm (such as ordered-subsets expectation maximization) is unrolled, or unfolded (8), into a deep network—the word *deep* meaning that there are many successive steps, as indeed in any piece of computer code. Iterative reconstruction is thus nothing more than a deep cascade of successive operations, each operation taking the raw PET data and progressively transforming it (by a series of operations, primarily forward and back projections) into a reconstruction of the PET radiotracer distribution. Deep learning is then integrated into the unfolded reconstruction to provide rich, data-informed, prior information to the iterative process, which makes repeated use of the actual raw data throughout. Thus, the benefits of

decades of reconstruction research are combined with the power of the AI paradigm (i.e., learning from high-quality reference datasets), allowing the manifold of feasible PET images to be used as a powerful, yet relatively safe (data-consistent), prior in the image reconstruction process. Compared with direct AI methods, the need for training data in these unrolled methods is reduced by orders of magnitude, as the physics and statistics of PET data acquisition do not need to be learned from scratch. Furthermore, their scope for generalization to unseen data is better than that of direct methods, as has been demonstrated in other imaging inverse problems (9).

The third main category of AI for PET reconstruction acts on existing standard reconstructed PET images. Such postprocessing is much simpler to implement, and this is where advances are being quickly made, with commercial options already available (such as subtlePET [<https://www.subtlemedical.com>], which seeks to map low-count [25% dose] PET images to their full-dose equivalents). Research in this area is burgeoning, with a myriad of differing deep network mappings being proposed, to denoise, upgrade, and even mimic state-of-the-art PET reconstructions from higher-count data (10).

At present, nearly all AI methods for PET reconstruction have leaned heavily on convolutional neural network (11) mappings. However, the surge of more advanced data-mixing architectures, such as the immense success of transformers (12), with their powerful self-attention mechanism for rapid learning of long-range contexts in data, has yet to reach the PET reconstruction community, but it is sure to come. These highly successful architectures should deliver still more powerful ways of harnessing all acquired PET data to generate feature-rich manifold embeddings, benefiting clinical imaging tasks and even ultimately aiding management of the patient pathway.

PROBLEMS TO TACKLE AND OUTLOOK

There have, however, been ongoing expressions of concern regarding AI. For example, in the context of MRI the risk of hallucinations, artificial features, and instability has been studied (13). Such problems, even evidenced in physics-informed approaches (unrolled iterative methods), will need comprehensive investigation, research, and resolution for PET image reconstruction in order to deliver the robustness required for clinical imaging.

A crucial part of such research will be the need for benchmark datasets through which new AI algorithms for PET image reconstruction can be assessed. Such datasets ideally need international collaboration and contributions from clinicians and researchers in reconstructions from multiple institutions. Such datasets have already existed for decades in the image processing community and have been established more recently in the deep learning, computer vision, and MRI communities (e.g., CIFAR, MNIST, ImageNet, and fastMRI (14)). Ideally, benchmark datasets for PET image reconstruction should be provided and linked with particular clinical tasks (e.g., neurological disorder diagnosis or tumor detection).

Furthermore, to have confidence in the high image quality that can be delivered by AI approaches to image reconstruction, the arrival of evidential deep learning is timely. Also known as Bayesian deep learning, these approaches not only would provide high-quality reconstructed PET images but also deliver unequivocal indications of the AI's uncertainty (known as epistemic uncertainty) in various regions and details of the image—information that would be crucial during clinical reading.

Although supervised learning remains central to current developments in PET reconstruction, the field will need to exploit larger datasets for which the costly ground truth labels or targets are not known. Unsupervised pretraining of networks has shown great potential in computer vision, and image reconstruction models could very likely benefit from pretraining with unlabeled data, followed by fine tuning with the labor-intensive supervised labels. Better still, self-supervised learning paradigms should prove useful. In essence, instead of providing explicit, labor-intensive example inputs and outputs, only example data are provided, along with instructions on how to create the set of inputs and targets from the data for supervised learning. Self-supervised approaches have enabled training of huge-scale language models, including powerful transformer-based architectures such as GPT-3.

CONCLUSION

AI is here to stay, and validated PET reconstruction that makes use of its power will deliver images of enhanced clinical benefit, compared with methods that ignore its capabilities. Yet to arrive at this point it will be necessary to build confidence, and 2 approaches may help. First, adoption may need to be in a gentle, progressive fashion. At the very simplest level, deep learning can provide optimization of merely the degree of standard image smoothing, which has low risk but also a reduced degree of benefit. This small step up from our existing regularized reconstruction methods could allow use of AI to decide how much anatomical (CT or MRI) guidance information can reliably be applied for PET reconstruction.

Second, to ensure safe adoption of more sophisticated AI methods, it may prove necessary to use routes such as evidential deep learning, in which, for example, epistemic uncertainty is clearly expressed alongside the images. The AI output would thus be twofold: “this is the best estimate of the image for the patient” and “this is my confidence level for each detail and region in the image.”

The methods that are set to flourish will harness all our knowledge of physics, math, and statistics for PET reconstruction and synergistically combine these with the learning power of AI, with feasible demands for training data. Simply put: there is no reason to learn from scratch that which we know well already, and conversely there is no reason to insist on simple mathematical expressions for complex images. For example, we cannot analytically derive or program what a feasible PET image should look like, but deep learning can do this with ease.

Finally, the endpoint assessment of the impact of AI reconstruction on clinical tasks, preferably with well-understood benchmark datasets, will of course be essential. Without question, in the development and validation of AI for reconstruction, critical feedback from clinicians will be needed more than ever.

DISCLOSURE

No potential conflict of interest relevant to this article was reported.

REFERENCES

1. Senior AW, Evans R, Jumper J, et al. Improved protein structure prediction using potentials from deep learning. *Nature*. 2020;577:706–710.
2. Salama MA, Hassanien AE, Mostafa A. The prediction of virus mutation using neural networks and rough set techniques. *EURASIP J Bioinform Syst Biol*. 2016; 2016:10.

3. Kermany DS, Goldbaum M, Cai W, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*. 2018;172:1122–1131.e9.
4. Mehranian A, Reader AJ. Deep learning for PET image reconstruction. *IEEE Trans Radiat Plasma Med Sci*. 2020;5:54–64.
5. Reader AJ, Corda G, Mehranian A, Costa-Luis Cd, Ellis S, Schnabel JA. Deep learning for PET image reconstruction. *IEEE Trans Radiat Plasma Med Sci*. 2021;5:1–25.
6. Zhu B, Liu JZ, Cauley SF, Rosen BR, Rosen MS. Image reconstruction by domain-transform manifold learning. *Nature*. 2018;555:487–492.
7. Häggström I, Schmidtlein CR, Campanella G, Fuchs TJ. DeepPET: a deep encoder-decoder network for directly solving the PET image reconstruction inverse problem. *Med Image Anal*. 2019;54:253–262.
8. Gregor K, LeCun Y. Learning fast approximations of sparse coding. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. Omnipress; 2010:399–406.
9. Gilton D, Ongie G, Willett R. Neumann networks for linear inverse problems in imaging. *IEEE Trans Comput Imaging*. 2020;6:328–343.
10. Schramm G, Rigue D, Vahle T, et al. Approximating anatomically-guided PET reconstruction in image space using a convolutional neural network. *Neuroimage*. 2021;224:117399.
11. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE*. 1998;86:2278–2323.
12. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. arXiv website. <https://arxiv.org/abs/1706.03762>. Published June 12, 2017. Revised December 6, 2017. Accessed July 27, 2021.
13. Antun V, Renna F, Poon C, Adcock B, Hansen AC. On instabilities of deep learning in image reconstruction and the potential costs of AI. *Proc Natl Acad Sci U S A*. 2020;117:30088–30095.
14. Knoll F, Murrell T, Sriram A, et al. Advancing machine learning for MR image reconstruction with an open competition: overview of the 2019 fastMRI challenge. *Magn Reson Med*. 2020;84:3054–3070.
15. AI in PET image reconstruction v1.2. figshare website. https://figshare.com/articles/figure/AI_in_PET_image_reconstruction_v1_0/14685915. Published May 27, 2021. Accessed July 27, 2021.