Journal of Nuclear Medicine, published on July 8, 2021 as doi:10.2967/jnumed.121.262303

AI for PET Image Reconstruction

Andrew J. Reader[1] and Georg Schramm[2]

[1]School of Biomedical Engineering and Imaging Sciences, King's College London, UK

[2]Department of Imaging and Pathology, Division of Nuclear Medicine, KU/UZ Leuven, Belgium

1. Introduction

Artificial intelligence (AI) continues to deliver remarkable impact in numerous andhighly diverse fields, from physics, natural language processing, finance, human resources to image processing, protein folding [1] and prediction of viral mutations[2]. In broad terms, AI is any technology which can *learn* how to perform tasks from example data or experiences. This 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 ground breaking technologies, can initially lead to concerns, scepticism and even ethicalquestions. In particular, use of AI in medical imaging has demonstrated immensepotential [3], but a key question is how much we can trust AI in the formation of images that inform clinical decisions, where lives of patients are often at stake.
This brief article will consider the methodologies, benefits and concerns regardingAI for the case of the formation, or *reconstruction*, of PET images [4], and will focus on a sub-discipline of AI, namely *deep learning*. We will define deep learning below, and use this term interchangeably with AI.

2. Understanding AI and deep learning

So what is deep learning exactly? Deep learning can be considered as a sequence of steps which operates on input data to perform a desired task, where the steps arelearned from example inputs and 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 upon experience and knowledge of imaging physics, maths and statistics, can nowalso be integrated into the learning AI paradigm. Better still, state of the art imagereconstruction methods can likely be made even more reliable with AI-informed refinement.

However, AI has been frequently misunderstood, either due to 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 partly originates from the highly successful use of deep learning in computer vision tasks,where 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 (CNNs) have resulted in large uptake and application of CNNs to other tasks, where 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-parameterised architectures has made it easy to dismiss any chance of understanding (let alone designing) these sophisticated non-linear mappings, fuelling AI scepticism. As for conventional mathematical perspectives on feasibility of optimisation and overfitting to limited data, 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 optimisation, regularisation and generalisation.

Hence the rapid progress of AI methods, sometimes with loss of principled designchoices and often to the surprise of conventional mathematical thinking, has resulted in concern over its interpretability and trustworthiness. This has not been helped by reduced levels of rigour arising from the surge of innovation and exciting successes. But block box concerns and conventional mathematical views on optimisation are becoming dated perspectives, particularly in the context of deep learning for signal and image processing, where 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).

3. Why use AI for PET image reconstruction?

In applying AI to image reconstruction for PET, we are recognising 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 and/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.

First, let's recall what PET image reconstruction actually is: it is the use of raw list-mode or projection data acquired from a PET scan in order 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 mm, and the temporal resolution is of the order of many seconds. These limitationsare 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, improving spatial and temporal resolution, through accurate modelling of the imaging physics and statistics, and through use of prior information (including from CT or MR). 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 poses ongoing challenges toPET image reconstruction.

This is where AI can make a huge difference, in two main ways. Firstly, with sufficient example data, AI can learn the vast (but nonetheless highly restricted) realistic set of PET images that can 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 discards considerable amounts of *a priori* information.In contrast, AI's learning of the manifold of all feasible PET images can be used tomake 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.

Secondly, 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 PETscan data into latent feature representations which best serve our desired goals.
This includes 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 encodekey 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 alsoits ability to learn how to *decode* from these latent representations, to generate outputs for various desired tasks. Thus 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.

4. How can we use AI in PET image reconstruction?

There are currently three main approaches to using AI in PET reconstruction. The first group of approaches, *direct AI* (e.g. AUTOMAP [5] or DeepPET [6]), 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 be easily understood by comparison to conventional curve fitting and regression tasks, except in the case of deep learningof PET reconstruction we are performing regressions with extremely high- dimensional vectors. The input raw PET data are fully 3D sets of measured (time- of-flight) sinograms (with ~10^8 - 10^9 bins), for mapping to output 3D images (with
~10^7 voxels). At present, these direct deep learning methods look to be impractical,having only been demonstrated for small 2D reconstructions (e.g. 128×128 images), as they have colossal demands for computational memory and training setsizes (>10^5 datasets). Furthermore, they may not generalise well for unseen data (i.e. scan data that are too far from the example training data). Early tests of direct methods for real data 2D PET reconstructions have delivered images which 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 OSEM) is “unrolled” or unfolded [7]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 to 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 generalisation to unseen data is better than direct methods, as has been demonstrated in other imaging inverse problems [8].

The third main category of AI for PET reconstruction acts on existing, standard reconstructed PET images. Such post processing is much simpler to implement, and this is where advances are being quickly made, with commercial options already available (such as subtlePET [9], 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 CNN [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, is still 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, benefitting clinical imaging tasks and even ultimately aiding management of the patient pathway.

5. Problems to tackle and outlook

There have however been ongoing expressions of concern regarding AI. For example, in the context of MR imaging the risk of hallucinations / artificial features has been studied, and the risk of instability [13]. Such problems, even evidenced in physics-informed approaches (unrolled iterative methods) will needcomprehensive 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 reconstruction from multiple institutions. Such datasets have already existed for decades in the image processing community, and more recentlyin the deep learning, computer vision and MRI communities (e.g. CIFAR, MNIST,ImageNet, fastMRI [14]). Ideally, benchmark datasets for PET image reconstruction should be provided and linked with particular clinical tasks (e.g. diagnoses of neurological disorders, or tumour detection).

Furthermore, to have confidence in the high image quality that can be delivered byAI approaches to image reconstruction, the arrival of *evidential deep learning* is timely. Also known as Bayesian deep learning, these approaches would not only 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, which would be crucial prior to clinical reading.

While 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/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 unlabelled data, followed by fine tuning with the labour-intensive supervised labels. Better still, self-supervised learning paradigms should prove useful. In essence, instead of providing explicit, labour-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 asGPT-3.

6. Conclusion

AI is here to stay, and validated PET reconstruction which makes use of its powerwill deliver images of enhanced clinical benefit, compared to methods that ignoreits capabilities. Yet to arrive at this point it will be necessary to build confidence, and two approaches may help. First, if adoption is to take place, it may need to bein a gentle, progressive fashion. At the very simplest level, deep learning can provide optimisation of merely the degree of standard image smoothing, with lowrisk, but at reduced capability of course. This could be a small step up from our existing regularised reconstruction methods, using AI to decide how much anatomical (CT or MRI) guidance information can be reliably used for PET reconstruction.

Secondly, to ensure safe adoption of more sophisticated AI methods, it may prove necessary to use routes such as evidential deep learning, where, 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 which are set to flourish will harness all our knowledge of physics, maths and statistics for PET reconstruction, and synergistically combine these withthe learning power of AI with feasible demands on training data. Simply put: *thereis no point learning from scratch that which we know well already,* and conversely*there is no point insisting 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 end-point assessment of 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.

References

1. Senior, A.W., et al., *Improved protein structure prediction using potentials from deeplearning.* Nature, 2020. **577**(7792): p. 706-710.
2. Salama, M.A., A.E. Hassanien, and A. Mostafa, *The prediction of virus mutation using neuralnetworks and rough set techniques.* EURASIP J Bioinform Syst Biol, 2016. **2016**(1): p. 10.
3. Kermany, D.S., et al., *Identifying Medical Diagnoses and Treatable Diseases by Image-BasedDeep Learning.* Cell, 2018. **172**(5): p. 1122-1131 e9.
4. Reader, A.J., et al., *Deep Learning for PET Image Reconstruction.* IEEE Transactions onRadiation and Plasma Medical Sciences, 2020.
5. Zhu, B., et al., *Image reconstruction by domain-transform manifold learning.* Nature, 2018. **555**(7697): p. 487-492.
6. Haggstrom, I., et al., *DeepPET: A deep encoder-decoder network for directly solving the PET image reconstruction inverse problem.* Med Image Anal, 2019. **54**: p. 253-262.
7. Gregor, K. and Y. LeCun, *Learning fast approximations of sparse coding.* Proceedings of the 27th International Conference on International Conference on Machine Learning, 2010: p. 399-406.
8. Gilton, D., G. Ongie, and R. Willett, *Neumann Networks for Linear Inverse Problems in Imaging.* IEEE Transactions on Computational Imaging, 2020. **6**: p. 328-343.
9. subtlePET, *https://www.subtlemedical.com*
10. Schramm, G., et al., *Approximating anatomically-guided PET reconstruction in image space using a convolutional neural network.* Neuroimage, 2021. **224**: p. 117399.
11. LeCun, Y., et al., *Gradient-based learning applied to document recognition.* Proc IEEE, 1998. **86**: p. 2278-2323.
12. Vaswani, A., et al., *Attention is all you need.* arXiv preprint arXiv:1706.03762, 2017.
13. Antun, V., et al., *On instabilities of deep learning in image reconstruction and the potential costs of AI.* Proc Natl Acad Sci U S A, 2020.
14. Knoll, F., et al., *Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge.* Magn Reson Med, 2020. **84**(6): p.3054-3070.

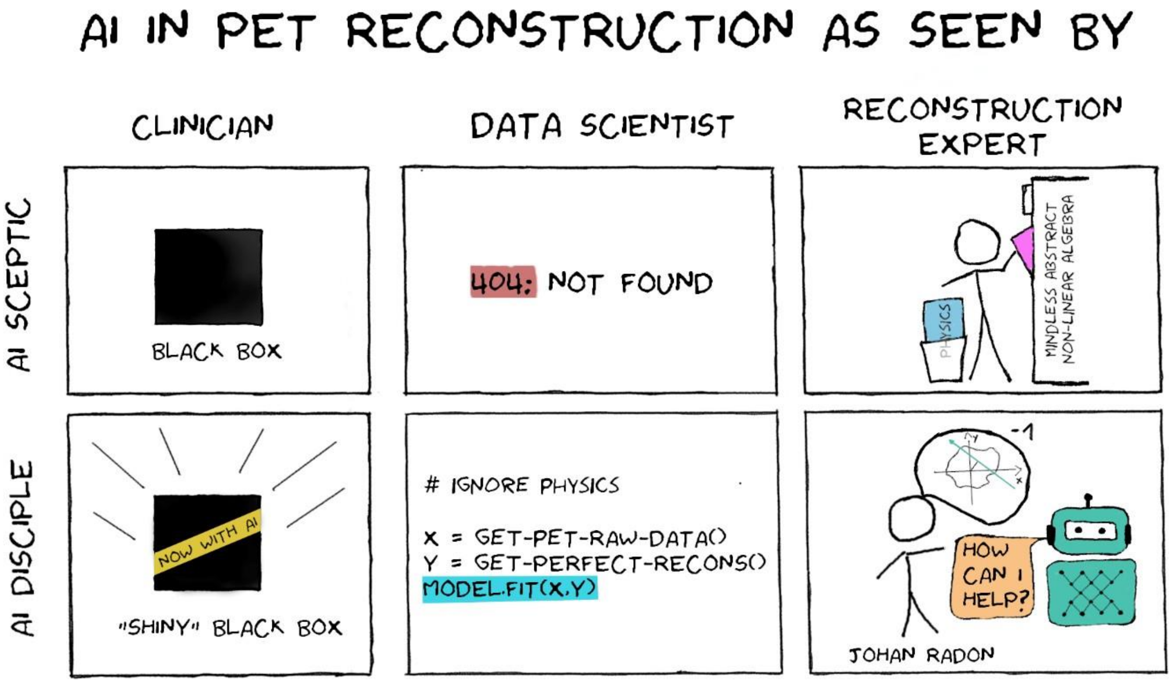

Figure 1: “AI in PET reconstruction as seen by”. https://doi.org/10.6084/m9.figshare.14685915.License CC BY 4.0. Georg Schramm and Andrew J. Reader 2021.