



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/152179/>

Version: Accepted Version

Article:

Gilbert, FJ, Smye, SW and Schönlieb, C-B (2020) Artificial intelligence in clinical imaging: a health system approach. *Clinical Radiology*, 75 (1). pp. 3-6. ISSN: 0009-9260

<https://doi.org/10.1016/j.crad.2019.09.122>

© 2019 The Royal College of Radiologists. Published by Elsevier Ltd. All rights reserved.
This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Clinical Radiology

Artificial intelligence in clinical imaging – a health system approach

--Manuscript Draft--

Manuscript Number:	
Full Title:	Artificial intelligence in clinical imaging – a health system approach
Article Type:	Editorial or Commentary
Corresponding Author:	Fiona J Gilbert University of Cambridge Cambridge, UNITED KINGDOM
Corresponding Author Secondary Information:	
Corresponding Author's Institution:	University of Cambridge
Corresponding Author's Secondary Institution:	
First Author:	Fiona J Gilbert, FRCR
First Author Secondary Information:	
Order of Authors:	Fiona J Gilbert, FRCR Stephen W Smye, PhD Carola Bibiane-Schönlieb, PhD
Order of Authors Secondary Information:	

Artificial intelligence in clinical imaging – a health system approach

Fiona J Gilbert MD¹, Stephen W Smye MA, MSc, PhD², Carola-Bibiane Schönlieb MMath, PhD³

¹ Department of Radiology, University of Cambridge School of Clinical Medicine
Box 218, Cambridge Biomedical Campus
Cambridge
CB2 0QQ
UK
Telephone: + 44 1223 336890 e-mail: fjg28@medschl.cam.ac.uk

² NIHR Clinical Research Network, School of Population Sciences and Health Services Research
Faculty of Life Sciences & Medicine
Kings College London, 6th Floor, Addison House, Guy's Campus,
London SE1 1UL
UK
Telephone: +44 0207 848 8370 e-mail: steve.smye@nhr.ac.uk

³ Department of Applied Mathematics and Theoretical Physics, [Centre for Mathematical Sciences](#),
Wilberforce Road, Cambridge CB3 0WA, UK
Telephone: +44 1223 765000 e-mail: cbs31@cam.ac.uk

Corresponding author;

Fiona J Gilbert, Department of Radiology, University of Cambridge School of Clinical Medicine
Box 218, Cambridge Biomedical Campus
Cambridge
CB2 0QQ
UK
Telephone: + 44 1223 336890 e-mail: fjg28@medschl.cam.ac.uk

Artificial intelligence in clinical imaging – creating a system

Manuscript type

Editorial Comment or Special Report

Key points

The development and application of Artificial Intelligence to Radiology requires an approach which encompasses a health system. The UK government and NHS are creating an ecosystem to facilitate academic/industrial partnerships aimed at accelerating the creation of relevant and robust AI tools which will improve the development and delivery of healthcare imaging. A series of recent initiatives are described which will drive the development and adoption of AI in clinical imaging.

The development of ever more sophisticated imaging technologies, such as computerised tomography (CT), Magnetic Resonance Imaging (MRI), Ultrasound and Positron Emission Tomography (PET), has led to a significant increase in the amount of data generated per patient[1]. At the same time, healthcare systems around the world are struggling to integrate and analyse this wealth of information, due to a shortage of trained radiologists and the sheer size and complexity of the datasets.

Academia and industry are focusing on developing artificial intelligence (AI) techniques – encompassing machine learning (ML) and deep learning (DL) in particular – for analysing, interpreting, categorising and annotating clinical images. Progress in AI imaging technology is being driven by the rapidly expanding processing power of GPUs, falling costs of computing and data storage, the availability of large datasets for training and significant financial input from private and commercial investors and government sources.

However, while AI has the potential to transform clinical imaging practice around the world by improving productivity and performance, there are significant issues that need to be resolved before imaging records are used at scale for training and AI is adopted in clinical practice. Although our focus is on the UK landscape, the challenges we face are relevant to the international research community of academic, clinical and industry partners working to speed the translation of AI imaging technologies into routine clinical practice.

The value and pitfalls of AI in imaging are increasingly well-rehearsed in the literature [2,3]. It is clear that for the potential of AI in imaging to be realised, a whole-systems approach is required, which supports ready access to well-curated, annotated, data-sets by a skilled multi-disciplinary workforce,

in a manner which is trusted by patients and the public and which enables the development and robust evaluation of novel analytical techniques.

Bringing the power of AI to bear on clinical imaging is a multidisciplinary effort, requiring close collaboration between academic researchers, clinicians, industry, government agencies, healthcare professionals and patients to develop solutions that are safe, effective and integrate into clinical workflows. Engineering these partnerships and creating a research ecosystem in which they can flourish will rely on strategic direction and investment from national bodies, including government-funded research organisations, industry and professional societies.

Creating an ecosystem

In the UK, three linked elements of an emerging ecosystem to deliver AI for imaging are of note: the recent Life Sciences Sector Deal [4] sets out an investment plan for the sector which includes investment in AI, to be delivered in partnership with industry and academia, by UK Research & Innovation (UKRI), aimed at placing the UK at the forefront of AI and data revolution; the National Institute for Health Research [5] - probably the world's largest integrated health research system, with more than £500m pa of research infrastructure embedded in the NHS designed to deliver high quality research at pace - is evolving in order to ensure the health system is well-placed to deliver high quality studies of AI in imaging; and NHSX [6], which brings together the Department of Health and Social Care, NHS England and NHS Improvement with oversight of NHS Digital to drive transformation of health services by digital technology, with a particular focus on setting standards and developing platforms, thereby providing a framework for the development and adoption of AI in health.

Several specific initiatives are now underway; funded by the industrial challenge strategy fund, UKRI ran a funding competition in 2018 for £50m (~\$65 million) to create a network of academic centres working with industry and focusing on developing AI in radiology and digital pathology for the NHS. A further £33 million (~\$42 million) in funding was leveraged from universities, charities and companies ranging from small start-ups to major multinational companies [7]. Based in Leeds, Oxford, Coventry, Glasgow and London – but each with partners across many parts of the UK – the centres will develop AI for radiology and digital pathology. Working with NIHR, the five funded academic/industrial partnerships have been charged with delivering AI tools for clinical testing within the NHS in a three year time frame.

While much of the attention in AI is focused on data and algorithm development, it is essential to remember the importance of skilled scientists and clinicians. This requires investment in multidisciplinary training programmes spanning the entire career pathway from studentships to fellowships and beyond, covering both clinical and non-clinical researchers. More should be done to encourage cross-disciplinary talent transfer, not just for medical doctors moving into AI research but also bringing physicists, mathematicians, engineers and computer scientists into closer proximity with biomedical researchers. This could extend to cross-disciplinary degrees, fellowships or even clinical placements to encourage people with AI skills to move into medical imaging. In February 2019, UKRI announced funding awards of £100m to create 16 new Centres for Doctoral Training (CDTs) based at 14 UK universities with 300 partners, including AstraZeneca, Google and Rolls-Royce, and NHS trusts [8]. Project partners are investing £78 million in cash or in-kind contributions and partner universities are committing a further £23 million, resulting in an overall investment of more than £200 million. Applying AI to healthcare, including imaging, is an important element of the programme. This CDT platform should provide numerous opportunities to promote the interdisciplinary exchange that is vital for AI in imaging to flourish. The NIHR is also working with

UKRI (Health Data Research UK) to develop a clinical academic development group in health data science which will help address some of these challenges.

One important issue to be considered is the financial incentives available. Given that AI-based companies can offer significantly higher salaries than those available for academic researchers, developing effective partnerships with industry will be an important means of developing and retaining talent as the field grows. In the UK, UKRI and the NIHR now have a range of fellowship programmes which encourage cross-disciplinary working and industry placements, whilst an NIHR-wide imaging programme also aims to support the development of a trained imaging research workforce in the NHS.

There is a risk that recruitment into training programmes will be hampered by misplaced concerns that radiology will become redundant thanks to the advent of AI technologies, further depressing these disciplines that are already struggling. Instead, we must focus on promoting the model of the 'centaur' – a highly trained human working together with an AI to achieve more than would be possible alone. The Royal College of Radiologists, responsible for training standards is actively engaged in the development of AI and is ensuring the workforce is appropriately informed.

Building a pipeline for development and validation

There are two main challenges to be overcome when bringing AI techniques into clinical imaging: development of the tools themselves and their subsequent clinical validation and approval. These two strands must run in parallel and be closely intertwined – there is no point developing an impressive algorithm if it cannot be integrated into day-to-day service delivery, demonstrate its effectiveness and utility in real life situations and meet the conditions for regulatory approval. Furthermore, any AI-based imaging system must fit seamlessly into established clinical workflows – for example, integrating into existing workstations rather than operating in a standalone unit –

otherwise it is unlikely to be widely adopted. It also has to demonstrate increased productivity, better patient outcomes and cost effectiveness, particularly in settings with stretched healthcare budgets. The developing role of NHSX, working in partnership with key partners, including the Royal College of Radiologists, is likely to be particularly important.

The most significant limiting factor in the development of AI technologies is the availability of sufficiently large, good quality training data. Keeping images in the AI pipeline is crucial for human interpretation and validation of resulting algorithms. Ideally, datasets should be uniformly acquired with standardised protocols across all sources, consistently annotated and anonymised or pseudo-anonymised, depending on where and how they will be used. Annotation currently requires human input, creating a bottleneck in the process due to the lack of trained radiologists and pathologists. This is driving the trend towards unsupervised learning techniques – where salient features are recognised without human intervention – as well as the use of computer-generated training data created through generative adversarial networks.

A recent national initiative funded by UKRI is particularly relevant; Health Data Research UK is a collaboration of 22 universities and research institutes, clustered into six regions designed to use health data securely to derive new knowledge and scientific discovery. The six new data centres will create a research infrastructure platform on which research with digital data can be streamlined and optimized.

There is also a shift away from using training data comprising processed images that are optimised for human viewing only and towards integrating raw acquisition data and physics models of the acquisition into the AI workflow. This would allow the creation of homogenised image datasets and their direct optimisation for diagnosis and treatment planning. There are also issues around

interoperability and regional variations; an algorithm that works with data generated on one machine may not perform as well with images gathered from another.

Despite the enthusiasm for developing AI-based imaging tools, rigorous clinical validation of these technologies remains a major challenge [3]. Unlike pharmaceutical companies, which must navigate a highly complex and well-established regulatory environment in order to gain approval for novel therapies, the regulatory framework for AI-based clinical technologies is still playing catch-up. Any validation test should be appropriate to the level of risk involved. For example, an algorithm designed to triage patients in a fracture clinic can tolerate more error than one designed to assess correct placement of a nasogastric feeding tube, where the outcomes of misplacement are life-threatening. Therefore, there is a critical need for the development of robustness measures and uncertainty quantification for AI techniques and their requirements in varying clinical settings.

Local differences in practice and patient populations also pose challenges; will a tool that has been developed using a population of breast cancer patients in Scotland be relevant to women in the southern states of the USA, or even in the south of England? Training and test datasets therefore need to be truly representative of the patient population to which the algorithm will be applied, or the specific patient population should be specified as part of the regulatory process. It is possible to imagine a solution for DL technologies where there is a core algorithm with 'add-ons' allowing for domain adaptation and consequently account for these local variations.

Replication and reproducibility are significant concerns for clinical validation, particularly for results produced by proprietary algorithms generated by commercial organisations that are reluctant to reveal their 'special sauce'. The benchmark for all AI imaging technologies should therefore be published, peer reviewed clinical trials, with as much transparency around the methodology, algorithm, training and test datasets and possible sources of bias as possible. There should also be

an accurate characterisation of failure cases: it is just as important to understand any 'blind spots' as it is to demonstrate an impressive detection rate. The need to develop robust methods for evaluating the clinical utility of AI algorithms in imaging is therefore clear and, working with partners, the NIHR is reviewing processes and novel approaches to evaluation which will increasingly characterise this important area.

As more AI-based platforms come to market, a concerted effort needs to be put into establishing standardised independent test datasets to demonstrate accuracy, sensitivity and specificity, analogous to the validation and quality control panels that are available for molecular diagnostics. These must be large enough to avoid the problem of 'learning to the test' or overfitting and allow for frequent retesting.

While it is important that any regulatory processes do not create an unnecessary barrier to clinical use, it is vital that there is adequate oversight to ensure that AI technologies are safe, effective and accepted by patients and the public. ML/DL software is likely to fall under the banner of 'medical devices' and will therefore be subject to having to gain CE accreditation in Europe or FDA approval in the US, which brings a requirement for post-marketing surveillance. It is also necessary to consider regular retesting and revalidation of AI-based algorithms. Several new AI tools have gained CE marks and FDA approval based on scant and often unpublished clinical data, and there are concerns that a failure to properly validate and monitor the application of these technologies in the real world could lead to potentially serious errors, risking the loss of public and professional trust.

The British Standards Institute (BSI, Medicines & Healthcare products Regulatory Agency (MHRA) and the Association for the Advancement of Medical Instrumentation (AAMI) have now published a joint report on Machine Learning in Healthcare summarising the key recommendations based on the

UK and US workshops held [9]. Further work, including participatory workshops, is planned aimed at developing approaches to regulation.

Data governance

Patient-derived data lies at the heart of any AI-based imaging system and is therefore subject to informed consent, national legal and regulatory frameworks, and societal norms. The introduction of the EU General Data Protection Regulation (GDPR) has clarified the requirements for organisations that gather and process personal data and, in our view, has been an enabler of research. However, there is still much confusion among the research community about how to navigate through the regulatory process – particularly for small commercial organisations – and further advice from regulatory bodies such as the UK Health Research Authority (HRA) would be welcome. There would also be value in the development of standardised national and international data-sharing agreements and contracts, which are already becoming common in pharmaceutical drug development and trials.

AI research in medical imaging would benefit from new models for data accessibility, moving from the idea of data sharing to one of data access. Several academic and commercial organisations have accumulated extremely large datasets that could be of great use to the research community, and we would all benefit from the development of platforms that allow researchers to come and use cleaned, curated data within an organisational firewall with the appropriate permissions. This would have the advantage of democratising data science, reducing barriers to entry for small organisations and countries with less investment in their health data infrastructure.

Public trust

Patients should be at the heart of research, not only as beneficiaries of these new technologies but also as partners and participants at all stages of the process from design to delivery – a principle that

lies at the heart of the NIHR. As well as being ethically correct, patient and public engagement and involvement makes research more effective, encouraging trial participation and retention and ensuring that the results of research are more likely to bring meaningful benefits. Despite recent high-profile scandals around the mis-use and leakage of personal data, the UK public remains broadly supportive of the use of patient data for medical research, even by commercial organisations [10]. However, levels of public trust are likely to vary by country and are currently being investigated through programmes such as the Wellcome Trust Global Monitor [11].

However, just as what was acceptable practice in medical research fifty years ago is looked upon with horror today, we should be mindful that attitudes can change over time. There is a growing public suspicion of large privately-owned technology companies that gather and control personal data, whose priorities ultimately lie with their owners or shareholders rather than patients and the public, and the AI research community should continue to actively engage with patients and the public to monitor their concerns. For example, it is currently accepted that patients should not receive financial reimbursement for the use of their data or a share of the profit from any commercial product derived from it. That may change with the advent of blockchain technologies allowing individuals to control access to their personal data or even monetise it, as we are starting to see in the field of genomics, which are likely to impact upon public trust [12].

Finally, academic and industry researchers should consider how to deliver effective communication about AI-based technologies to patients, public and health professionals. It may not be necessary to completely explain the 'black box' of each algorithm – after all, we do not expect doctors or patients to know the precise biological mode of action of every drug – but efforts need to be made to show how these tools work, the data that they are derived from, and their benefits and limitations. Building a culture of transparency and public understanding around the use of AI in medical imaging will help to secure trust and confidence in this exciting field as it moves into the future.

References

1. Pelka O, Nensa F, Friedrich CM, Annotation of enhanced radiographs for medical image retrieval with deep convolutional neural networks. PLoS ONE 2018 13(11): e0206229. doi.org/10.1371/journal.pone.0206229
2. Lambin P, Leijenaar RTH, Deist TM, Peerlings J, de Jong EEC, van Timmeren J, Sanduleanu S, Larue RTHM, Even AJG, Jochems A, van Wijk Y, Woodruff H, van Soest J, Lustberg T, Roelofs E, van Elmpt W, Dekker A, Mottaghy FM, Wildberger JE, Walsh S. Radiomics: the bridge between medical imaging and personalized medicine Nature Reviews Clinical Oncology 2017 14, 749–762 doi:10.1038/nrclinonc.2017.141
3. Hosny A, Parmar C, Quackenbush J, Schwartz LH, Aerts HJWL. Artificial Intelligence in Radiology Nature Reviews Cancer 2018 18, 500-510, doi:10.1038/s41568-018-0016-5
4. <https://www.gov.uk/government/publications/life-sciences-sector-deal> accessed 26 June 2019
5. <https://www.nihr.ac.uk/> accessed 26 June 2019
6. <https://www.nhsx.nhs.uk/> accessed 26 June 2019
7. <https://www.gov.uk/government/news/artificial-intelligence-to-help-save-lives-at-five-new-technology-centres> accessed 26 June 2019
8. <https://www.ukri.org/news/200m-to-create-a-new-generation-of-artificial-intelligence-leaders/> accessed 26 June 2019
9. <https://www.bsigroup.com/globalassets/localfiles/en-gb/about-bsi/nsb/innovation/mhra-ai-paper-2019.pdf> accessed 27 February 2019
10. <https://wellcome.ac.uk/sites/default/files/public-attitudes-to-commercial-access-to-health-data-wellcome-mar16.pdf> accessed November 19, 2018

11. <https://wellcome.ac.uk/what-we-do/our-work/public-views-science-and-health> accessed November 20, 2018
12. Huang EC, Pu C, Chou YJ, Huang N. Public Trust in Physicians-Health Care Commodification as a Possible Deteriorating Factor: Cross-sectional Analysis of 23 Countries. *Inquiry*. 2018;55:46958018759174.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Funding

Funding from the National Institute for Health Research (NIHR), the Cambridge Centre for Mathematical Imaging in Healthcare, and the Engineering and Physical Sciences Research Council is gratefully acknowledged. FG is a consultant for Google DeepMind; She receives research funding from GE Healthcare

This editorial was developed following a cross-disciplinary workshop in November 2018, hosted by the Cambridge Centre for Mathematical Imaging in Healthcare, the NIHR Imaging group and the Newton Gateway to Mathematics at The Alan Turing Institute in London, UK. We are grateful for the insights provided by the invited speakers listed below and all the participants. Our special thanks go to Ben Glocker and Charlene Liew for discussions on the editorial and Kat Arney at First Create the Media Ltd for writing services. Presentations from the workshop are available to view online at <https://gateway.newton.ac.uk/event/tgmw64/programme>