Radiology: Artificial Intelligence

Artificial intelligence in clinical imaging – a health system approach

Journal:	*Radiology: Artificial Intelligence*
Manuscript ID	Draft
Manuscript Type:	Special Report
Manuscript Categorization Terms:	Computer Applications-General (Informatics) < 6. TOPICS, Unsupervised learning < Type of machine learning < 8. MACHINE LEARNING, Social implications < ELSI < Application Domain < 8. MACHINE LEARNING

SCHOLARONE™
Manuscripts

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Abbreviated title page

Artificial intelligence in clinical imaging – development and delivery

Abstract

The development and application of Artificial Intelligence to Radiology requires an approach which encompasses a health system. Key elements of this approach, which are of global relevance, are described in the light of recent initiatives by UK government. A network of academic centres focusing on developing AI in radiology and digital pathology has been created with industrial partnerships within the NHS and the research infrastructure of the UK National Institute of Health Research (NIHR). Close collaboration is required between academic researchers, clinicians, industry, government agencies, healthcare professionals and patients to develop solutions that are safe, effective and integrate into clinical workflows with multidisciplinary training to ensure an upskilled workforce. The NHS has large archives of processed digital images acquired in relatively uniform standards but optimised homogeneous annotated datasets are relatively infrequent. However by using larger quantities of images the AI algorithms can be adapted to this more "messy" data and is arguably more useful as this reflects real life radiological practice. Rigorous clinical validation of these technologies requires thoughtful approaches as some AI tools triaging workflow are low risk compared to those creating a diagnostic decision support application. The EU General Data Protection Regulation (GDPR) has clarified the requirements for organisations that gather and process personal data and the AI field would benefit from international data sharing agreements. Transparency and building public trust is important for ensuring acceptability of implementation of these AI tools into our healthcare systems.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

The development of ever more sophisticated imaging technologies, such as computerised tomography (CT), different forms of Magnetic Resonance Imaging (MRI), Positron Emission Tomography (PET) and digital pathology, 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 pathologists and the sheer size and complexity of the datasets available.

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.

For example, in November 2018, the UK Government announced funding of just under £50 million (~$65 million) for a network of academic centres focusing on developing AI in radiology and digital pathology. A further £33 million (~$42 million) in funding has been leveraged from universities, charities and companies ranging from small start-ups to major commercial players. Furthermore, the UK National Institute of Health Research (NIHR) is 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, in partnership with patients, industry, charities and UK Research & Innovation (UKRI). Working with the devolved nations, the NIHR is evolving in order to ensure the UK health system is well-placed to deliver high quality studies of AI in imaging.

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 it is adopted in clinical practice at scale. 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.

Creating an ecosystem

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. The need to pivot towards ever-increasing automation and AI technology in healthcare is particularly pressing in countries like Singapore, where an ageing population and restrictions on non-native workers are putting pressure on services.

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.

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.

There is a risk that recruitment into training programmes will be hampered by misplaced concerns that radiology and pathology 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. We should remember the lessons learned from the introduction of spreadsheets in the late 1970s: fears that this new software would completely remove the need for accountants couldn't have turned out to be further from the truth. Instead, by removing the drudgery of basic tasks, the technology revolutionised the profession by freeing up human brains to focus on tasks for which they are better equipped.

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 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.

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 make

of 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 [2]. 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.

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 sufficient 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.

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 [3]. 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 [4].

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 [5].

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.

Acknowledgements

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

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. https://wellcome.ac.uk/sites/default/files/public-attitudes-to-commercial-access-to-health-data-wellcome-mar16.pdf accessed November 19, 2018

4. https://wellcome.ac.uk/what-we-do/our-work/public-views-science-and-health accessed November 20, 2018

5. 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.