The Royal College of Radiologists
Response to:
House of Lords Select Committee on Artificial Intelligence – 6 September 2017

Introduction

1. The Royal College of Radiologists (RCR) works with our 10,000 members to improve the standards of practice in the specialties of clinical radiology and clinical oncology.

2.The Royal College of Radiologists welcomes the opportunity to give evidence to this inquiry as clinical radiology and clinical oncology are at the forefront of the use of Artificial Intelligence (AI) in a healthcare setting. For the purposes of this response the RCR is using the definition of AI as the branch of computer science dealing with the simulation of intelligent behaviour in computers. This response focuses on the short and medium term effects of AI.

Summary

- The members and fellows of the RCR are at the forefront of the development of AI in healthcare. We believe that AI may be part of the solution to the current workforce crises in both clinical radiology and clinical oncology but it cannot replace human clinicians.
- The development of AI in healthcare will require a huge amount of patient data. The NHS has these data available but it is crucial all data used are fully anonymised at source (before leaving hospital firewalls), and that patients understand how and why their data are used. The NHS must also get the best deal for any data used by private companies.
- The government has a crucial role to play in the development of AI in healthcare. This includes ensuring there is investment in infrastructure and workforce. Developing ways of sharing the patient data needed to develop AI and ensuring there is a regulatory system that protects both patients and doctors.
- The RCR suggests that AI companies are urged to concentrate upon developing machine learning that can recognise normal plain X-rays of every part of the body with over 99.5% accuracy (e.g the wrist, the spine, the pelvis etc.). If these could be safely and reliably reported by AI it would be very helpful in assisting with the reporting of normal plain X-rays. This may free up radiologists to report all the abnormal imaging studies and to perform hands-on interventional work and may alleviate some of the workforce pressures on radiology departments. Clinical judgement will still be essential for assessing individual cases. For example a patient who has a negative X-ray may still have significant problem that requires a radiologist to assess what further investigation is needed.
- For AI companies to develop software to recognise normal X-rays from all parts of the body, they need to train their machine learning using huge amounts of normal plain X-ray data. These data already exist on the PACS (picture archive and communication system) archives in every NHS hospital in the UK, but to safeguard patient privacy they need to be released in a completely anonymised form to a central, independently regulated national database. The RCR is developing a specification stipulating how these normal X-ray data can be anonymised at source before passing through the firewalls of hospitals. Funding is required for the RCR (as an independent expert professional body) to trial and test a prototype of this process, so that a very large national database of normal X-ray studies can safely and quickly be built up for the AI companies to use, for a fee.
- To ensure only quality plain X-ray data (true normals) are contributed to this national database archive, NHS radiologists would need to check, by double reporting, that all X-ray studies submitted are normal. Funding would need to be provided for this.

Evidence

The pace of technological change

1. **What is the current state of artificial intelligence and what factors have contributed to this? How is it likely to develop over the next 5, 10 and 20 years? What factors, technical or societal, will accelerate or hinder this development?**

3. AI is already in use throughout radiology and oncology and both specialties are likely to be heavily affected by AI over the next 20 years. The RCR does not believe that AI will be able to replace human clinical radiologists or clinical oncologists as AI will not have explanatory power, so cannot investigate cause. AI will also not replace a doctor's judgement, creativity or empathy.

4. The NHS handles over 40 million imaging investigations per year[1] and AI has the capability of developing sophisticated scheduling algorithms that dynamically match the needs to the patients, and to the capacity of the resources. All modern radiological imaging is already acquired, viewed, reported and archived electronically, and all modern radiotherapy (excepting skin cancer) is planned on electronic images, using complex suites of computer programmes. For those reasons, AI will make stronger inroads in these areas of healthcare than in others.

5. For interventional radiology, a subspeciality of radiology, AI may assist with deciding the pathway for each patient (treat, not treat or formal review) for procedures like stroke thrombectomy and angioplasties.

6. Clinical oncology was one of the first specialities routinely to computerise large elements of its work, initially for radiotherapy planning. Modern radiotherapy planning requires the routine use of complex software, often running on dedicated hardware, and is used to produce optimal radiotherapy plans.

7. AI will also have huge potential for breast imaging. There are approximately two million women screened per annum in UK[2] and images are read at a rate of 55 per hour. Replacing one of two breast screen readers with AI (all screening mammograms are double read by a radiologist, advanced practitioner or breast physician) would free up a significant amount of time for a service under severe stress due to staffing shortages. Many consultants are now due to retire at the same time, 30 years after the breast screening programme was set up.[3]

8. As AI develops it is likely that more complex clinical scenarios will be able to be analysed and the interplay between many comorbidities will be understood. However when benefits of therapy are minimal (as in palliative treatment towards the end of life) patients may prefer to explore these value options with a clinician. Key to decision making are the needs and wishes of patients and their families, which may not always align with the outcome of an algorithm.

9. Technical limitations will be one of the biggest factors which may hinder development including computer power and the availability of data. There needs to be investment in computers with parallel processing capability and NHS IT will need to be upgraded to cope with increasing demands. Another limitation will be the need for huge amounts of anonymised patient data for training AI algorithms; we have considered this further in question four. Another limitation will be the availability of anonymised patient data for machine learning. The RCR is working on a specification outlining how

[1] NHS England, Diagnostic Imaging Dataset Annual Statistical Release 2014/15 https://www.england.nhs.uk/statistics/wp-content/uploads/sites/2/2014/11/Annual-Statistical-Release-2014-15-DID-PDF-1.1MB.pdf
[2] Health and Social Care Information Centre, Breast Screening Programme, England Statistics for 2014-15 http://content.digital.nhs.uk/catalogue/PUB20018
[3] Royal College of Radiologists, The breast imaging and diagnostic workforce in the United Kingdom https://www.rcr.ac.uk/system/files/publication/field_publication_files/bfcr162_bsbr_survey.pdf

plain X-rays can be anonymised before being used by AI companies and as an independent body we could oversee the development of a national database.

2. Is the current level of excitement which surrounds artificial intelligence warranted?

10. There has been some unwarranted hype over AI in the past but there is now a conjunction of software, hardware and data that make further development of AI possible. However, there remains a risk of over-hyping AI in particular areas. In particular, the adoption of AI into clinical practice faces several barriers. Widespread adoption or spread must be done in a measured way that ensures that there is sufficient research, rigorous clinical testing, clinical training and regulation. AI companies should be encouraged to focus on developing software that can recognise normal plain X-rays rather than anything more complex as that is where the most value can be added,

Impact on society

3. How can the general public best be prepared for more widespread use of artificial intelligence?

11. The public can be prepared for the use of AI in healthcare by fully addressing issues around consent and ownership of data, and by transparent scientific demonstration of the accuracy and reliability of AI in domains where it is to be used. AI systems work best with the very large amount of data available at NHS level rather than at Trust level. Access to these data requires the general public to understand what their data are being used for and have confidence that the development of AI and use of their data are well regulated. Public confidence is likely to depend upon ensuring that the uses of data clearly benefit other patients, and that data are not released to commercial parties without clear safeguards.

12.As AI begins to play a role in clinical radiology and clinical oncology it will be crucial that information is made available about the role that AI has in patient care, and crucially what role AI cannot play, and so will continue to be done by humans. There is a risk that current rhetoric around AI may lead people to believe erroneously that machines can entirely replace humans in this area.

13. The general public can also be prepared for widespread use of AI by the development of education to support the future workforce. In particular higher education needs to ensure that students have the skills to utilise AI techniques. A major challenge at present in the real world applications of AI, such as in healthcare, is bridging the understanding gap between domain experts (clinicians) and AI experts.

4. Who in society is gaining the most from the development and use of artificial intelligence and data? Who is gaining the least? How can potential disparities be mitigated?

14. All levels of society that use the NHS stand to benefit from AI, given that imaging is indispensable for the whole of modern medicine and surgery, including cancer pathways and trauma. There is potential for large swathes of routine work to be done by AI which may be disruptive to staff. Within imaging, the workforce of highly skilled and flexible workers will be advantaged as AI has the potential, by "weeding out" and accurately reporting all normal plain X-rays, to free up capacity to allow radiologists and oncologists to concentrate on the more complex studies and work. This approach across imaging will enable radiologists to make use of their high level skills for those patients who need them.

Public perception

5. Should efforts be made to improve the public's understanding of, and engagement with, artificial intelligence? If so, how?

15. It is essential that the public understand how AI can benefit them, particularly in the context of the delivery of healthcare. The adoption and implementation of AI may improve outcomes, and free up time so that doctors can spend more time with patients and the more complex aspects of their jobs. If the results of certain investigative tests can reliably and accurately be provided by AI, this would mean many patients would wait less time for their test results. The use of AI in such circumstances may make up for workforce vacancies, and can be used day and night 365 days of the year.

16.However to develop AI in clinical radiology and clinical oncology vast amounts of patient data are needed to train the AI software, which means the public must be understand that their data will be safely and thoroughly anonymised, and know how their data will be used. Whilst lots of the early AI work is done in academic institutions, ultimately academic partnerships with industry are required for commercialisation. Patients and the public may well be wary of their data being used for commercial benefit and at the moment this is a grey area.

17. The government should protect intellectual property generated from patient data with careful academic/industry contracts which reward the NHS for its contribution. This could include shares in the profits from commercial products which can be fed back into the NHS/academia, or negotiating free installation and access to the AI products for NHS patients and staff, as well as payment for access to these anonymised data. The concerns about sharing patient data could be reduced by normalising data sharing nationally in the UK e.g. all imaging studies to be used for AI research and training to be rigorously anonymised at source (within hospital firewalls), and then transferred into a national repository, overseen by an independent body such as the RCR. Ideally there would be the implementation of an opt-out system for patients, even when data have been rigorously anonymised. Most importantly there must be transparent systems in place to reassure the public.

Industry

6 What are the key sectors that stand to benefit from the development and use of artificial intelligence? Which sectors do not?

18. Medical imaging is perfectly placed to benefit from advances in AI. It is one of the few areas that has large data pools of high quality curated data (in the form of imaging studies and their reports) thus overcoming one of the main hurdles in AI development. The government needs to invest now in medical imaging research if the UK is to play a significant role in AI in the future. The AI technology industries need for data, especially labelled data, puts the NHS in a unique position.

19. The NHS should be mindful that the data AI models rely upon are key. Therefore, negotiations between NHS organisations and AI developers should be carefully monitored to avoid handing over data for little, or no, reward. At present, the tendency is to release the data to developers. However, this approach is not without risk so there should be the development of centralised AI resources available within the NHS. Such an NHS AI Institute could be relatively small, and yet provide shared resources and expertise across many domains. Patients may feel more comfortable with their data being used within the NHS rather than outside and have the process of anonymising data overseen by the appropriate body, in the case of medical imaging this should be the Royal College of Radiologists

20. The risk with AI is that a small number of industry leaders, once ahead of their competitors, will develop an unassailable advantage, due to the nature of iterative reinforcement learning. For example, an imaging product becomes useful and is adopted. New images are passed through the product as it is being used and as a result the product improves as it 'sees' more examples. Given the global nature of the AI industry, the UK must invest now in order to maintain a foothold in this industry in the future. There is potential for the NHS to achieve substantial strides in medical AI given that it is the world's largest single-payer healthcare provider, and continues to enjoy public support.

7. How can the data-based monopolies of some large corporations, and the 'winner-takes-all' economies associated with them, be addressed? How can data be managed and safeguarded to ensure it contributes to the public good and a well-functioning economy?

21. Any transfer of patient data between the NHS and private companies must get value for money for the NHS. The NHS has invested significant resource in developing banks of data that are now being used by technology companies. It is vital that this effort is recognised and appropriately rewarded by private companies who will profit from it. Given the complexity of the area, and the potential for network effects, there should be some consideration to the idea that an NHS AI Institute should supervise such deals. The risk is that otherwise commercial entities may be able to cherry-pick deals with individual NHS Trusts.

Ethics

8. What are the ethical implications of the development and use of artificial intelligence? How can any negative implications be resolved?

22. There are ethical issues around consent for data to be used by AI and also consent of patients for AI to be used in their treatment. We have explored some of these in issues in questions three and five above. If the data are thoroughly anonymised before passing through hospital firewalls, this should provide reassurance to patients

23. Private companies using NHS data must ensure that they abide by the same ethical standards as the NHS for studies. This includes ensuring applications to ethics committees for use of data are completed. Security of patient data is also important; one way of overcoming this could be for data to be held by the NHS in a centralised resource. However, there are wider ethical issues around trust and liability in AI-enabled healthcare. There remain open questions about when failure to use an AI system would be negligent, or how errors in a combined human/ AI system would be addressed.

9. In what situations is a relative lack of transparency in artificial intelligence systems (so called 'black boxing') acceptable? When should it not be permissible?

24. AI has the capacity to process information to degrees of complexity that far exceeds the comprehension of the human mind. Much work is being done to decipher 'black box' algorithms as this enables improvements to be made, and this will shed some light on the process but will not be comprehensive. Indeed test-validation of complex algorithms may be unfeasible in practice. More generally, however, there is a gap between the validation of algorithms and the validation of their implementation clinically. Although there are guidelines on software engineering (e.g. as used in the Space Shuttle), these are rarely followed in current AI software. This should be considered in regulation.

The role of the Government

10 What role should the Government take in the development and use of artificial intelligence in the United Kingdom? Should artificial intelligence be regulated? If so, how?

Data and infrastructure

25. For AI to be a success there needs to be upfront investment in IT systems. Currently effective IT is not available and this will slow the implementation of AI. Such investment will mean improved decision making which will have a cost saving.

26. NHS data should be viewed as a key strategic resource, modern AI is extremely data dependent, and a licensing model that allows collaboration with commercial entities while ensuring that benefits

are returned to the NHS, is vital. The nature of AI-based systems, and their dependence on large amounts of data, creates the possibility of creating a 'winner-takes all' situation where once a system has established a lead, it becomes impossible for competitors to overcome its advantages. The organisation of such systems also has the potential to reinforce economic inequality, with large gains flowing to relatively small groups of people/companies. It seems unlikely that using large amounts of NHS patient data to generate profits for small numbers of companies would be acceptable to the public.

27. There should be consideration of developing some centralised, sharable resource of AI computing, especially for building models over large NHS datasets. Such infrastructure could be virtual but it would provide a central point of contact and expertise to provide overlap between the clinical and AI worlds, with appropriate safeguards. The RCR can assist in ensuring that data used is fully anonymised.

Regulation and governance

28. The governance of the use of AI in healthcare is unclear, and problematic. At present, clinicians are held responsible for deciding to use, and interpret the results of most AI tools. Such tools include auto-contouring software in radiotherapy planning or decision support software for breast cancer chemotherapy.

29. However, as AI becomes more autonomous, then governance becomes more difficult. It is not clear at what point failure to use an AI system would become negligent. Linked to this is the relatively poor development and documentation process for much software development. Although there are methodologies for writing safety-critical software, these are currently rarely used in AI systems. In addition, the tendency for software to be provided as a service, rather than as a downloadable product, makes it easy for the provider incrementally to upgrade the product. While this has many advantages, it may mean that the AI service is not stable over time. Auditing the results of such a service is therefore challenging.

30. Without a clear governance framework for managing the risk associated with using AI tools, their introduction is likely to be slow, and haphazard. This framework would need to include multiple elements, including the data used to build such tools, software engineering practices and updates, and a clear process for understanding and assigning risk. This discussion will need to involve a range of partners, including regulators and indemnity organisations

31. Regulation of AI in healthcare will be crucial: not just the regulation of software itself by Medicines and Healthcare products Regulatory Agency (MHRA) but the impact AI may have on the regulation of healthcare professionals. If a doctor uses a treatment recommended by an algorithm where do the professional responsibilities of the doctor end? A regulatory framework should consider such issues.

Liability

32. Legal liability is often stated as a major societal hurdle to overcome before widespread adoption of AI becomes a reality. If mistakes are made in the course of treatment or diagnoses where AI is used there will be a need for robust legislation to cover liability.

Investment

33. Given the amount of AI talent, both commercially and academically, and the presence of large datasets (especially healthcare-related) there is a potential for the UK to be a world leader in AI. The government needs to ensure it seizes this opportunity by ensuing that there is adequate investment and support for AI. This should include providing funding for independent bodies like the RCR to oversee the development of systems of anonymise patient data.

34. The government needs to clarify immediately the status of foreign workers as the elite in this field are few in number, and the UK needs to ensure that it is a welcoming place to work. The barriers in this sector, real, or perceived, are damaging to the workforce

35. The National Institute for Health Research (NIHR) should consider setting up an NIHR AI Bioresource similar to the approach taken to genomics, with the possibility of it being spun out into a private company equivalent to Genomic England.

Learning from others

11. What lessons can be learnt from other countries or international organisations (e.g. the European Union, the World Economic Forum) in their policy approach to artificial intelligence?

36. The American College of Radiology has established a Data Science Institute to guide the appropriate development and implementation of AI tools to help radiologists improve medical imaging care.

Dr Nicola Strickland
President
Royal College of Radiologists
September 2017