

Artificial Intelligence in Radiology: An Exciting Future, but Ethically Complex



Author:	*Adrian P. Brady 1. Mercy University Hospital, Cork, Ireland 2. University College Cork, Ireland 3. Second Vice President, European Society of Radiology (ESR), Vienna, Austria *Correspondence to adrianbrady@me.com
Disclosure:	The author has declared no conflicts of interest.
Received:	19.01.21
Accepted:	25.02.21
Keywords:	Artificial intelligence (AI), machine learning, medical ethics, radiology, radiomics.
Citation:	EMJ Radiol. 2021;2[1]:54-57.

BACKGROUND

“Let me start by saying a few things that seem obvious. I think if you work as a radiologist, you’re like the coyote that’s already over the edge of the cliff but hasn’t yet looked down, so doesn’t know there’s no ground underneath him. People should stop training radiologists now. It’s just completely obvious that within 5 years, deep learning is going to do better than radiologists, because it’s going to be able to get a lot more experience. It might be 10 years, but we’ve got plenty of radiologists already. I said this at a hospital, and it didn’t go down too well.”¹

With those words at a 2016 Creative Destruction Lab (CDL) seminar on ‘Machine Learning and the Market for Intelligence’ in Toronto, Canada, Dr Geoff Hinton provided radiologists the world over with an uncomfortable prediction of their obsolescence (and provided a piece of video that always gets attention from audiences during speeches about artificial intelligence [AI] and radiology). Dr Hinton, an English/Canadian cognitive psychologist and computer scientist, is, fittingly, the great-great-grandson of George Boole.

There have been many other such predictions in recent years, some from sources that know less about the subject than Dr Hinton. In October

2020, the Dutch Finance Minister, Wopke Hoekstra, said: “The work of the radiologist to a significant extent has become redundant, because [...] a machine can read the images better than humans who studied 10 years for it.” He also commented that the same changes were occurring with supermarket checkout operators.²

Whatever one thinks about the value of such apocalyptic prognostications for the demise of the specialty (or about the lack of understanding of the work that underpins them), it is true that the advent of AI tools will change (indeed, is already changing) radiology practice. The era of spending long periods painstakingly perusing hundreds of images to identify tiny lung nodules on a chest CT will soon be past, and unlamented. AI will automate many tedious and imperfect aspects of radiologists’ work and allow for the focus of more time and effort on higher-level cognitive tasks (such as deciding which of the many lung nodules identified by the AI tool are significant and what they mean) and collaborative input into diagnosis and management of patients, the components of radiologists’ jobs that many glib commentators do not recognise or understand.

CAN WE DEVELOP ARTIFICIAL INTELLIGENCE ETHICALLY?

So, can we assume that AI in radiology will be a boon to radiologists, to patients, and to society as a whole? That depends on who develops and integrates it, how they do it, and their motivations.

AI developments in all disciplines are led by innovative, intelligent, and dedicated researchers or software developers. These people need support and encouragement. They also need practical financial resources and these are frequently supplied by large data companies, directly or indirectly. The rapid movement of 'big tech' players into healthcare AI indicates the importance they attach to these developments and, not incidentally, the monetisable value they foresee arising from access to and use of healthcare data. Developing AI tools that perform accurately and effectively requires access to large amounts of verified 'ground truth' data, which can be used to train and validate algorithms. In radiology this usually involves large numbers of labelled imaging studies, upon which the machine learning algorithm practises and hones its functions. Obtaining the necessary imaging data and labelling its contents can be a very resource-intensive task; the huge resources of big data companies can be a perfect match for this need. However, the means by which the data are obtained, manipulated, and used must be controlled to avoid inadvertent or deliberate ethical missteps. Under the provisions of the General Data Protection Regulation, all European Union (EU) citizens own and control their own sensitive, personal, and identifiable data. If patients' imaging data are to be made available to AI developers, patients' privacy and data ownership rights must be protected, their consent to use of their data must be solicited and secured, and they must agree to any possible financial ramifications of their information being used to develop potentially profitable outputs. They also need to be sure that their personal information is secure (or ideally irreversibly removed from the data) and will not be used by data companies to target them in ways that have nothing to do with the stated use for which they have consented.

Moving beyond the issues of access to and anonymisation of imaging data, another ethical

issue that arises in AI in radiology is data bias. This is usually inadvertent, arising from an algorithm having been trained on data that doesn't accurately represent the population on which the tool is ultimately deployed, whether because of under- or over-representation of particular population subsets in training data or fundamental differences between populations. Examples abound of AI tools performing outstandingly in training but failing when applied to populations with different demographics or characteristics.³ Frequently, these difficulties cannot be identified in advance but their possibility should always be considered and anticipated when developing AI programmes.

SHOULD ARTIFICIAL INTELLIGENCE BE TRANSPARENT AND UNDERSTANDABLE?

The very nature of machine learning makes it difficult, if not impossible, for humans to follow and understand every step that takes place during the process whereby an AI algorithm arrives at an outcome (i.e., a 'decision'). This has been referred to as the 'black box' situation, whereby inputs are provided to the algorithm, an output is delivered, but we cannot follow how that output was arrived at. Interpretability, explainability, and transparency seem like attractive attributes for an AI algorithm if its developers want to engender public trust, representing, respectively, the ability to understand the workings of the AI model, to explain in understandable terms what happens when the model makes a decision, and the capability of an outside observer to visualise and understand what happens within the model. Unfortunately, these are not necessarily always viable or desirable goals. If we can follow and understand every step of an AI process, we surrender some of the potential benefits of machine learning, including the capacity of the algorithm to achieve goals that are beyond the conscious mind. Furthermore, the more transparent an AI model is, the more subject it is to malicious attack and the less secure is its intellectual property, potentially reducing its commercial value to developers. There are ethical trade-offs inherent in developing and marketing useful AI products, balancing the need for understanding of the people on whom the product will ultimately be used, against the

advantages of harnessing computing power to augment human capability.

Public acceptance of the use of AI in medical care should not be taken for granted. Most would not be comfortable to accept without question that decision-making about their healthcare should be devolved to algorithms, without human oversight. Public knowledge of the role AI may play in future life is, as yet, underdeveloped. Sixty-five percent of American adults have been shown to be uncomfortable about delegating the task of making a medical diagnosis to a computer with AI.⁴ When asked about autonomous vehicles (AV; self-driving cars), in general, the public approve of AV that would sacrifice passengers for the greater good if faced with the choice of either running over pedestrians or sacrificing the occupants, and would like others to buy them. However, most people would themselves prefer to travel in AV that protect passengers at all costs.⁵ This level of inconsistency augurs poorly for public understanding and acceptance of AI in healthcare.

ETHICAL DANGERS IN ARTIFICIAL INTELLIGENCE

At a practice level, AI tools offer potential for misuse. It is not difficult to imagine a healthcare provider adapting an algorithm to drive medical decisions which will increase utilisation and profit, rather than be solely based on patient welfare. Equally, better-off patient groups could derive advantages over other subsets of the population from AI resources, which they can afford to access (the term 'liberal eugenics' has been used to describe this scenario). To some extent, these dangers also exist in conventional healthcare. Nonetheless, we should guard against increasing the potential for their occurrence.

Other practice-based ethical dangers of AI also exist. If something goes wrong after AI use in medicine, who is liable? Is it the doctor who used the tool, the institution that bought it, or the developer who brought it to market? Doctors' involvement in AI model development is desirable and beneficial; after all, we understand best how these tools may be applied in patient care. But this involvement opens up the possibility

of decisions about which AI tools to use, and how to use them, being exploited for personal commercial gain.

AI is, at heart, a mathematical function. As such, it can perform very well in classification tasks. It's less well-suited to more abstract concepts, such as determination of fairness, equality, and context. Writing code to embed the ability to weigh up the consequences of a decision or management recommendation for a patient or a patient's family, with all the associated calculations and choices that can depend on very individual circumstances, is difficult. Simultaneously ensuring that AI-supported decision-making always follows ethical principles adds another layer of complexity.

CONCLUSION

This paper attempts to outline some, but by no means all, of the ethical issues that arise when AI algorithms are in development and use in clinical practice. None of these issues is insoluble. Equally, none of them is simple or easy to resolve. Many of these potential problems can be lost sight of in the excitement of developing and implementing new tools, which have the potential to greatly benefit patients and to change medical practice for the better. We can best guard against deliberate or inadvertent unethical actions by educating all involved in AI about the moral risks that exist in the use of AI. To this end, in 2019, a joint group representing the European Society of Radiology (ESR), the American College of Radiology (ACR), the Radiological Society of North America (RSNA), the Canadian Association of Radiologists (CAR), the European Society of Medical Imaging Informatics (EuSoMII), the Society for Imaging Informatics in Medicine (SIIM), and the American Association of Physicists in Medicine (AAPM) published a comprehensive statement on the ethics of AI in radiology.^{6,7} This statement tries to explain the issues in detail. We do not yet have all the solutions to these ethical considerations, but acknowledging and understanding the problem is the first, necessary step to ultimately get the implementation of AI in radiology right, using its power for the benefit of individual patients and society, without harm or bias.

References

1. Creative Destruction Lab. Geoff Hinton: on radiology. 2016. Available at: <https://www.youtube.com/watch?v=2HMPRXstSvQ>. Last accessed: 4 March 2021.
2. Turner J, Ward P. Dutch debate intensifies over future shape of AI. 2020. Available at: <https://www.auntminnieeurope.com/index.aspx?sec=sup&sub=aic&pag=dis&ItemID=619384>. Last accessed: 20 October 2020.
3. Khullar D. A.I. could worsen health disparities. 2019. Available at: <https://www.nytimes.com/2019/01/31/opinion/ai-bias-healthcare.html>. Last accessed: 4 March 2021.
4. Morning Consult. National Tracking Poll 170401. 2017. Available at: https://morningconsult.com/wp-content/uploads/2017/04/170401_crosstabs_Brands_v3_AG.pdf. Last accessed: 4 March 2021
5. Bonnefon JF et al. The social dilemma of autonomous vehicles. *Science*. 2016;352(6293):1573-6.
6. Geis JR et al. Ethics of artificial intelligence in radiology: summary of the joint European and North American multisociety statement. *Insights into Imaging*. 2019;10:101.
7. Geis JR et al. Ethics of AI in radiology: European and North American multisociety statement. 2019. Available at: https://static-content.springer.com/esm/art%3A10.1186%2Fs13244-019-0785-8/MediaObjects/13244_2019_785_MOESM1_ESM.pdf. Last accessed: 4 March 2021

FOR REPRINT QUERIES PLEASE CONTACT: INFO@EMJREVIEWS.COM