



# The Challenges of Regulating Artificial Intelligence in Healthcare

## Comment on “Clinical Decision Support and New Regulatory Frameworks for Medical Devices: Are We Ready for It? - A Viewpoint Paper”

Martin McKee<sup>1\*</sup>, Olivier J. Wouters<sup>2</sup>

### Abstract

Regulation of health technologies must be rigorous, instilling trust among both healthcare providers and patients. This is especially important for the control and supervision of the growing use of artificial intelligence in healthcare. In this commentary on the accompanying piece by Van Laere and colleagues, we set out the scope for applying artificial intelligence in the healthcare sector and outline five key challenges that regulators face in dealing with these modern-day technologies. Addressing these challenges will not be easy. While artificial intelligence applications in healthcare have already made rapid progress and benefitted patients, these applications clearly hold even more potential for future developments. Yet it is vital that the regulatory environment keep up with this fast-evolving space of healthcare in order to anticipate and, to the extent possible, prevent the risks that may arise.

**Keywords:** Regulation, Clinical Decision Support, Artificial Intelligence

**Copyright:** © 2023 The Author(s); Published by Kerman University of Medical Sciences. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

**Citation:** McKee M, Wouters OJ. The challenges of regulating artificial intelligence in healthcare: Comment on “Clinical decision support and new regulatory frameworks for medical devices: Are we ready for it? - a viewpoint paper.” *Int J Health Policy Manag.* 2023;12:7261. doi:10.34172/ijhpm.2022.7261

### Article History:

Received: 18 March 2022  
Accepted: 7 September 2022  
ePublished: 19 September 2022

### \*Correspondence to:

Martin McKee  
Email:  
[martin.mckee@ishtm.ac.uk](mailto:martin.mckee@ishtm.ac.uk)

### Regulating Health Products

Regulation of health technologies must be rigorous, instilling trust among both healthcare providers and patients. Regulatory mechanisms have developed over time, with advances often following revelations of weaknesses in the regulatory process, such as those that allowed the teratogenic drug thalidomide to be prescribed in the 1960s.<sup>1</sup> The principles underlying the regulation of pharmaceuticals have been extended to other medical technologies. Now, as described in the accompanying paper by Van Laere and colleagues about the use of clinical decision support systems, regulators are working out how to deal with applications using artificial intelligence in healthcare.<sup>2</sup>

There are many challenges in regulating healthcare technologies. An obvious example is how to deal with the emergence of side effects not identified in the initial trials, either because they are rare, only develop after some time, or are only found in patients with characteristics not included in those trials. Other issues relate to biological agents where superficially minor differences in manufacturing processes can impact on safety and effectiveness<sup>3</sup> and challenges in recruiting enough patients for trials of drugs treating rare conditions.<sup>4</sup> There are, arguably, even greater challenges with medical devices. Regulators may differ in what they see as falling within their remit, which has resulted in weak control and supervision of medical devices in many settings. The

performance of the device may also vary according to the skill and experience of the operator.

Yet these hurdles are relatively minor compared to those involved in regulating the growing use of artificial intelligence in healthcare. As Van Laere and colleagues conclude in their viewpoint article, “designing a regulatory framework that achieves the right balance between promoting innovation and fast market access on one side and ensuring safety and quality on the other side is very challenging.” We agree, and in this commentary seek to complement their analysis by looking in more detail at some of the issues that arise in the regulation of artificial intelligence in healthcare.

### Artificial Intelligence in the Healthcare Sector

First, it may be useful to set out the scope for applying artificial intelligence in the healthcare sector, as some readers may be unfamiliar with its key characteristics. In essence, it seeks to improve on the decision-making process undertaken by human operators as they synthesise and interpret information and make decisions. Recent advances have incorporated probabilistic reasoning to deal with uncertainty and machine learning, whereby the algorithms improve with experience. Machine learning now underpins most of modern artificial intelligence and can be unsupervised, seeking a pattern in the data presented to it, or supervised, whereby it learns from information fed into it by a human who has labelled it (for

example, by adding the definitive diagnosis to a package of clinical data). Artificial intelligence can use a wide range of data inputs,<sup>5</sup> although most of the early applications in healthcare relied on visual images, such as those used in radiology (eg, positron emission tomography scans), pathology (eg, images of cells and tissues), or ophthalmology (eg, retinal pictures).<sup>6</sup> In due course this has expanded to the analysis of more complex three-dimensional images, such as those obtained at colonoscopy, and a wide range of physiological data, such as that generated by echocardiography, most often with the aim of making a diagnosis. In some cases, data are being linked in imaginative ways. For example, analysts have successfully combined data on symptoms with recordings of coughs to accurately diagnose respiratory infections,<sup>7</sup> and linked radiographic and longitudinal clinical data to offer a prognosis and inform subsequent monitoring and treatment.<sup>8</sup>

In these ways, artificial intelligence can not only improve the quality of care but, crucially, in health systems that are often constrained by numbers of skilled health professionals, also support increased activity. There is enormous potential to take advantage of the vast quantities of data that can now be collected on people engaged in everyday activities through wearable technology, such as the activity trackers contained within smartphones or devices that continuously monitor, for example, blood glucose levels.

As with medicines and medical devices, it is not possible to make generalisations about the performance of clinical decision support software that rely on artificial intelligence, but there is now considerable evidence that, in certain circumstances, some can perform as well as, or even better, than human decisionmakers.<sup>9-11</sup> Yet artificial intelligence is not a panacea, as recent experiences during the coronavirus disease 2019 (COVID-19) pandemic have shown.<sup>12</sup> In 2007 Weiner and colleagues used the term “e-iatrogenesis” to denote patient harm resulting from information technology.<sup>13</sup> Cabitza and colleagues have identified four broad risks.<sup>14</sup> These are that artificial intelligence may deskill health workers, whose performance may be degraded if the product is unavailable or dysfunctional; it may fail to take account of context, such as differences in patient mix in different settings; it may fail to take account of uncertainty, for example in categorising input data that are subject to inter-observer variability; and problems may arise from the opacity of the process. Burrell<sup>15</sup> identifies three aspects to this opacity. Two of these aspects, namely corporate secrecy by the provider and technical illiteracy by the user, can be overcome, at least in theory. But the third, the intrinsic complexity of the algorithm, cannot easily be addressed. Grote and Berens<sup>16</sup> illustrate the problem with reference to the common situation where two expert clinicians disagree. They can discuss the reasons for their disagreement but where a clinician and a machine disagree the conversation will be one-sided.

Finally, artificial intelligence is contributing to healthcare in other ways too. Biosimulation, in which the behaviour of chemical entities is analysed *in silico*, is becoming increasingly important in drug development.<sup>17</sup> In transcriptomics, which is the study of messenger RNA to ascertain which of an organism's genes are active, artificial intelligence is being

used to analyse genomic and transcriptomic data from microorganisms to detect antimicrobial resistance.<sup>18</sup>

### Challenges Facing Regulators

These developments have potentially profound consequences for clinical practice, but they also raise very difficult issues for regulators who are charged with protecting the public from unsafe and ineffective tools. We can identify at least five.

First, an artificial intelligence application, where utilised, is only one part of a complex clinical system. It will require data to be inputted in an acceptable form. But what if the data input device is inadequately calibrated, or the application has been trained on high-quality images but is presented with low-quality ones? Will the regulatory process be able to take this into account?

Second, the process of training the application may incorporate existing values and biases without making them explicit. For example, one study found that a white patient given a certain score by an application designed to estimate risk in primary care patients in the United States was deemed healthier than a black patient with the identical score.<sup>19</sup> This was because the outcome variable was based in part on cost of treatment, with black patients typically receiving less expensive care. Biases in algorithms might be reduced by granting analysts access to larger, more representative datasets, but that would mean combining data from different providers into a single application. Regulation that enables this, while putting safeguards in place to ensure it is done safely, ethically, and in a manner that maintains individual privacy, could go a long way in improving artificial intelligence systems in healthcare.

Third, where the application includes machine learning, its performance will change over time. This suggests that regulatory approval should be time limited. But how frequently should it be redone, given the trade-off between risk (which may be exceedingly difficult to estimate) and regulatory burden (which is measurable)? The US Food and Drug Administration has proposed a life-cycle process, from pre-market development to post-market performance, but how this will work out in practice is unclear.<sup>20</sup>

Fourth, artificial intelligence in healthcare can conflict with data protection legislation, which in many settings (such as those covered by the European Data Protection Regulation) requires that only data required for the purpose intended should be collected. Yet artificial intelligence applications are extremely data hungry and it is often very difficult, if not impossible, to determine what information is necessary for the algorithms to function and what is not.<sup>21</sup> It also raises issues of potential fraud: A recent World Health Organization (WHO) report has highlighted this danger, noting how a survey distributed by Facebook that was purported to be a psychological test was used to develop algorithms later used to influence elections.<sup>22</sup> In the context of healthcare, should private health insurers be able to secure access to sensitive information that could help them predict the risk of individuals requiring healthcare, then they could use these data (which they are not supposed to have access to) to illegally adjust premiums. This could impact millions of individuals if done at scale. While this was always a potential concern with any

illegal access to medical records, the opportunities created by artificial intelligence are immense.

Fifth, applications are gathering vast quantities of data, raising issues of privacy. It is possible to identify characteristics of the patient that they do not want to be recorded in their data. It is known that artificial intelligence can predict parameters such as chronological age from even quite limited radiological data, something that is not especially surprising.<sup>23</sup> However, Gichoya and colleagues have shown that deep learning algorithms can predict race with a high level of accuracy from a wide variety of radiological images.<sup>24</sup>

These are some of the main issues facing regulators assessing artificial intelligence as a diagnostic aid. However, there is one other area that, although in its infancy, should not be overlooked. Earlier we mentioned its use in drug design. Yet, as was the case with other technologies, such as nuclear energy, things can be used for both good and evil. While algorithms used in this way are typically designed to screen out toxicity, a group of researchers turned this on its head. In a proof of concept study they showed that, within a few hours, they could design analogues of known chemical weapons predicted to be even more toxic.<sup>25</sup> They call for greater awareness of the scope for dual use of artificial intelligence, ethics training for those involved, and channels for reporting potential abuses. In summary, this provides another, previously largely ignored challenge for those regulating artificial intelligence that falls on the margins of existing technology assessment models.

Given these issues, it will often be difficult to decide who is accountable if things go wrong. When a medicine is approved, the approval comes with conditions. These include the indications for use (ie, the condition that the product is used for) and perhaps patient characteristics such as age or renal function which should be considered when administering a product, or other medications with which it should not be given because of known interactions. The physician can still use it if these conditions are not met, as an off-label prescription, but then takes responsibility. Of course, even with correct use, a medication administered may still be unsafe, perhaps because there was a risk that should have been identified during development but was not or it was not stored in the right conditions. In such cases, the responsibility is clear. The situation is much more difficult with artificial intelligence-based applications. Is it the designer of the initial algorithms, the person responsible for entering the data such as the echocardiograph operator, or the clinician proposing treatment who must decide how much weight to place on the answer given by the software when it conflicts with other evidence visible to the clinician but not captured by the algorithm that is responsible?

### What Can or Should Regulators Do?

Van Laere and colleagues have described in detail the mechanisms that US and European authorities have put in place to regulate artificial-intelligence tools for clinical decision support.<sup>2</sup> They outlined difficulties in defining which products should be subject to such regulation and classifying the risk profiles of products. In an accompanying commentary, Maresova expanded on the regulation around

medical devices in the European Union.<sup>26</sup> There is hope that the European Union's proposed Artificial Intelligence Act, which would be the first comprehensive regulatory scheme for artificial intelligence worldwide and would divide products into risk categories, will usher in a new era in the regulation of artificial intelligence and establish a global standard for regulators and manufacturers.<sup>27</sup>

We agree with Van Laere and colleagues that neither set of regulators have yet clarified all of the issues that arise. Among the concerns that they raise, we see two as being especially challenging. The first, which has also been raised by the US Food and Drug Administration, is the need to devise a regulatory process that spans the entire life-cycle of the application.<sup>28</sup> This should be one that fosters innovation while ensuring patient safety, a difficult balance to achieve. It will require development of standards at all points on this cycle. Starting with the premarket authorisation, this would include the creation of a process similar to the phases of clinical trials undertaken with pharmaceuticals, although adapted to the different context. For example, this might include standards for validation of the algorithms, although given the many ways in which artificial intelligence can be used, we should not underestimate the challenges of doing this. While, superficially, it would seem that the ability to explain the logic would be desirable, as the WHO report noted, this risks inhibiting innovation.<sup>22</sup> However, there is an argument for including a requirement that algorithms should be capable of being evaluated independently. It would also include a test of patient benefit that would screen out applications that are really only data harvesting tools. Once the application has been introduced, there should be clear rules for when changes to the software were of sufficient importance to justify a further review to ensure that new risks had not been introduced. Finally, there could be a requirement for regular audits to be undertaken, at prespecified intervals, to identify potentially hazardous drift from the initial performance.

The second is one that is common to the regulation of conventional medical products, including pharmaceuticals. This is how the evaluations that inform the regulatory process may be conducted in populations that are unrepresentative of those on whom they will ultimately be used. For example, it is well known that many clinical trials of medicines exclude older people or those with multi-morbidity.<sup>29</sup> This is even more important when evaluating artificial intelligence applications. Furthermore, it is especially important to anticipate, investigate, and prevent algorithms from replicating or reinforcing existing biases. This depends on a high level of awareness of the risks but can be mitigated by measures such as the creation of training datasets that have been evaluated as having low risk of bias.<sup>30</sup>

None of this will be easy. While artificial intelligence applications in healthcare have already made rapid progress, these applications clearly hold even more potential for future developments. The innovations that they have already provided are bringing benefits to patients. However, the regulatory environment needs to keep up with this fast-evolving space of healthcare in order to anticipate and, to the extent possible, prevent the risks that may arise.

**Ethical issues**

Not applicable.

**Competing interests**

Authors declare that they have no competing interests.

**Authors' contributions**

Both authors contributed equally to this work.

**Authors' affiliations**

<sup>1</sup>Faculty of Public Health and Policy, London School of Hygiene & Tropical Medicine, London, UK. <sup>2</sup>Department of Health Policy, London School of Economics, London, UK.

**References**

- Eichler HG, Abadie E, Baker M, Rasi G. Fifty years after thalidomide; what role for drug regulators? *Br J Clin Pharmacol*. 2012;74(5):731-733. doi:10.1111/j.1365-2125.2012.04255.x
- Van Laere S, Muylle KM, Cornu P. Clinical decision support and new regulatory frameworks for medical devices: are we ready for it? - A viewpoint paper. *Int J Health Policy Manag*. 2022;11(12):3159-3163. doi:10.34172/ijhpm.2021.144
- Kirchhoff CF, Wang XM, Conlon HD, Anderson S, Ryan AM, Bose A. Biosimilars: Key regulatory considerations and similarity assessment tools. *Biotechnol Bioeng*. 2017;114(12):2696-2705. doi:10.1002/bit.26438
- Tudur Smith C, Williamson PR, Beresford MW. Methodology of clinical trials for rare diseases. *Best Pract Res Clin Rheumatol*. 2014;28(2):247-262. doi:10.1016/j.berh.2014.03.004
- Rajpurkar P, Chen E, Banerjee O, Topol EJ. AI in health and medicine. *Nat Med*. 2022;28(1):31-38. doi:10.1038/s41591-021-01614-0
- Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016;316(22):2402-2410. doi:10.1001/jama.2016.17216
- Porter P, Abeyratne U, Swarnkar V, et al. A prospective multicentre study testing the diagnostic accuracy of an automated cough sound centred analytic system for the identification of common respiratory disorders in children. *Respir Res*. 2019;20(1):81. doi:10.1186/s12931-019-1046-6
- Kehl KL, Elmarakeby H, Nishino M, et al. Assessment of deep natural language processing in ascertaining oncologic outcomes from radiology reports. *JAMA Oncol*. 2019;5(10):1421-1429. doi:10.1001/jamaoncol.2019.1800
- Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017;542(7639):115-118. doi:10.1038/nature21056
- Walsh CG, Ribeiro JD, Franklin JC. Predicting risk of suicide attempts over time through machine learning. *Clin Psychol Sci*. 2017;5(3):457-469. doi:10.1177/2167702617691560
- Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7
- Heaven WD. Hundreds of AI tools have been built to catch covid. None of them helped. 2021. <https://www.technologyreview.com/2021/07/30/1030329/machine-learning-ai-failed-covid-hospital-diagnosis-pandemic/>. Accessed May 20, 2022.
- Weiner JP, Kfuri T, Chan K, Fowles JB. "e-Iatrogenesis": the most critical unintended consequence of CPOE and other HIT. *J Am Med Inform Assoc*. 2007;14(3):387-388. doi:10.1197/jamia.M2338
- Cabitz F, Rasoini R, Gensini GF. Unintended consequences of machine learning in medicine. *JAMA*. 2017;318(6):517-518. doi:10.1001/jama.2017.7797
- Burrell J. How the machine 'thinks': understanding opacity in machine learning algorithms. *Big Data Soc*. 2016;3(1):2053951715622512. doi:10.1177/2053951715622512
- Grote T, Berens P. On the ethics of algorithmic decision-making in healthcare. *J Med Ethics*. 2020;46(3):205-211. doi:10.1136/medethics-2019-105586
- Maharao N, Antontsev V, Wright M, Varshney J. Entering the era of computationally driven drug development. *Drug Metab Rev*. 2020;52(2):283-298. doi:10.1080/03602532.2020.1726944
- Bhattacharyya RP, Bandyopadhyay N, Ma P, et al. Simultaneous detection of genotype and phenotype enables rapid and accurate antibiotic susceptibility determination. *Nat Med*. 2019;25(12):1858-1864. doi:10.1038/s41591-019-0650-9
- Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447-453. doi:10.1126/science.aax2342
- Food and Drug Administration. Artificial Intelligence and Machine Learning in Software as a Medical Device. 2021. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>. Accessed March 17, 2022.
- van Kolschooten H. EU regulation of artificial intelligence: challenges for patients' rights. *Common Mark Law Rev*. 2022;59(1):81-112. doi:10.54648/cola2022005
- World Health Organization (WHO). *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance*. Geneva: WHO; 2021.
- Eng DK, Khandwala NB, Long J, et al. Artificial intelligence algorithm improves radiologist performance in skeletal age assessment: a prospective multicenter randomized controlled trial. *Radiology*. 2021;301(3):692-699. doi:10.1148/radiol.2021204021
- Gichoya JW, Banerjee I, Bhimireddy AR, et al. AI recognition of patient race in medical imaging: a modelling study. *Lancet Digit Health*. 2022;4(6):e406-e414. doi:10.1016/s2589-7500(22)00063-2
- Urbina F, Lentzos F, Invernizzi C, Ekins S. Dual use of artificial-intelligence-powered drug discovery. *Nat Mach Intell*. 2022;4(3):189-191. doi:10.1038/s42256-022-00465-9
- Maresova P. Impact of regulatory changes on innovations in the medical device industry: Comment on "clinical decision support and new regulatory frameworks for medical devices: are we ready for it? - A viewpoint paper." *Int J Health Policy Manag*. 2022. doi:10.34172/ijhpm.2022.7262
- Heikkilä M. A quick guide to the most important AI law you've never heard of. 2022. <https://www.technologyreview.com/2022/05/13/1052223/guide-ai-act-europe/>. Accessed August 21, 2022.
- Hwang TJ, Kesselheim AS, Vokinger KN. Lifecycle regulation of artificial intelligence- and machine learning-based software devices in medicine. *JAMA*. 2019;322(23):2285-2286. doi:10.1001/jama.2019.16842
- Britton A, McKee M, Black N, McPherson K, Sanderson C, Bain C. Threats to applicability of randomised trials: exclusions and selective participation. *J Health Serv Res Policy*. 1999;4(2):112-121. doi:10.1177/135581969900400210
- Vayena E, Blasimme A, Cohen IG. Machine learning in medicine: addressing ethical challenges. *PLoS Med*. 2018;15(11):e1002689. doi:10.1371/journal.pmed.1002689