https://doi.org/10.1038/s41591-022-02165-8

How to regulate evolving AI health algorithms

By David W. Bates

Artificial intelligence algorithms have had mixed success in health, in part because regulation prevents them from evolving at the necessary rate.

Regulation of healthcare information technology has long been a thorny problem. Unlike drugs and devices that are relatively static, information technology such as apps, telehealth and artificial intelligence (AI) needs to evolve rapidly. The Food and Drug Administration (FDA) has the main authority to regulate these products in the USA, to assure that they are safe and effective, but its approval cycle typically takes months to years. This does not fit well with software, which needs to be updated weekly or monthly.

There is great excitement about the potential of AI to improve healthcare in a variety of ways, but such algorithms evolve rapidly and sometimes learn continuously. An approved AI algorithm can be locked but this has a disadvantage, as it may perform worse in sites outside of where it was trained, and a locked model that cannot learn will not improve.

Widespread use of AI in health will require interfaces that work with underlying clinical systems and tools so that predictions are sent to the appropriate clinician in real time, and accurate predictions that have both high positive predictive values and reasonable levels of sensitivity and specificity. These critical features will likely require models that can learn continuously.

This is not just a theoretical problem. An evaluation of one widely used algorithm, the Epic sepsis model, showed that the model performed poorly in many hospitals around the USA, which nonetheless continued to use it without appearing to notice how bad the performance was. It identified only 183 of 2,552 patients with sepsis (7%) who did not get timely antibiotics. It also failed to find around two-thirds of patients with sepsis, despite alerting on 18% of hospitalized patients.

A well-chronicled evaluation of a different sepsis algorithm found a modestly lower mortality rate, less organ failure and shorter length of stay in the intervention group, perhaps because the group used better technique and adhered more closely to the guidelines for algorithm development.

This varying efficacy suggests that how models are developed matters, which should inform new regulatory approaches. The FDA has recently issued guidance that it is considering "a total product lifecycle-based regulatory framework for these technologies that would allow for modifications to be made from real-world learning and adaptation, while ensuring that the safety and effectiveness of the software as a medical device are maintained". This would be a welcome departure from prior approaches. The FDA specifically called out sepsis as one of the use cases that it considers most important, but there will be many others.

A related issue is that AI always performs best in the dataset in which it is developed, and less well in future datasets, but not all models are validated in an external dataset. Moreover, there is a risk that commercial organizations will overstate the performance of their models and will run many AI models simultaneously, which will dramatically increase the scale of the challenge. Healthcare organizations may not want to get all AI models from one supplier, as important use cases will vary by institution, and there are hundreds or even thousands that could be beneficial.

Regulation of AI algorithms will need to be radically nimbler than prior approaches used to regulate drugs and devices. The FDA has previously regulated software used to collect, test, prepare, store and transport blood and blood products, with the resulting software being safe but with problematic usability. Regulation by the FDA should focus on safety and effectiveness of these applications. Achieving this will require the FDA to develop some new skills, as well as a deeper understanding of how different types of AI work and the extent to which it can be made explainable, and to streamline some of its internal processes, but the current commissioner is very familiar with the issues involved. The Health Information Technology and Economic and Clinical Health Act of 2009 aimed to spur such advances, but this has not yet been realized.

Better regulation would enable the healthcare industry to start using AI in many areas and to have confidence that it will improve the quality and safety of care, while reducing costs. It would also broaden the market for healthcare information technology companies, with the ultimate goal of better patient care.

David W. Bates[1,2,3] ✉

[1]Division of General Internal Medicine and Primary Care, Brigham and Women's Hospital, Boston, MA, USA. [2]Harvard Medical School, Boston, MA, USA. [3]Department of Health Care Policy, Harvard Medical School, Boston, MA, USA.

✉e-mail: dbates@bwh.harvard.edu

Published online: 24 January 2023

Competing interests

D.W.B. reports grants and personal fees from EarlySense, personal fees from CDI Negev, equity from ValeraHealth, equity from Clew, equity from MDClone, personal fees and equity from AESOP, and grants from IBM Watson Health, outside the submitted work.