Editorial

How to support the transition to AI-powered healthcare

Check for updates

To make health systems more sustainable in the long-term, incentivize artificial intelligence (AI) and digital technologies that are grounded on careful testing and real-world validation.

In a not-so-distant future, the world population will look much older than it used to be. A recent WHO report[1] estimated that 1 in 6 people will be over 60 by 2030. Healthcare systems in high-income countries are already experiencing the strain from shifting demographics, as a shrinking working age population needs to attend to an ever-growing aging population. Although life expectancy has been increasing steadily in the past 20 years in most countries, healthy life expectancy has not grown at the same pace.

Against this backdrop, governments and other societal forces are directing their attention to digital and computational tools that can reduce the cost of current healthcare systems without compromising standards of care, possibly even increasing their reach and quality. Recent studies[2-4] have offered a glimpse of what the future could look like, showing that computer vision algorithms can act effectively as additional 'eyes' in breast cancer screening, increasing the accuracy of case detection.

We are confident that this transition to 'AI-powered' healthcare will occur and that it has the potential to bring widespread public good. At the same time, we believe that these benefits will realize more steadily and more quickly with carefully designed clinical studies and evidence-based implementation of AI algorithms and devices in the real world. We are eager to support researchers and clinicians in this endeavor and clinical implementation will continue to be one of our top priorities for digital medicine content.

The transformative potential of AI does not come without risks. The questions of how AI interventions should be evaluated or when an intervention is ready for primetime are still open. In this respect, the fact that regulators are struggling to keep up with the pace of technological innovations in this field does not help. At present, digital and computational tools still hover in the grey area of medical devices, for which prospective clinical evaluation is often not required. Fears about the harmful use of AI, particularly the introduction of algorithmic biases that could skew or prevent someone from receiving appropriate care are real and could be catastrophic if scaled up. Preventable setbacks like these would only slow down the adoption of AI tools in the clinic and ultimately increase its costs in the long run. So what is the way forward to realize the immense potential of AI?

First, prospective testing and validation are crucial. There is extensive evidence that AI models have generalizability issues, meaning that an AI tool trained on one dataset may not offer accurate predictions when exposed to new data. For instance, a methodological review of the hundreds of machine learning models developed for COVID-19 screening[5] throughout the pandemic revealed that a large majority of them were problematic due to insufficient sample sizes, absent external validation and inappropriate performance evaluation. AI models also tend to have very different performance across population subgroups, typically favoring the majority one, for which they have seen the most data. That may result in worse outcomes for underrepresented groups. But even a so-called perfect model needs to be tested in its intended setting, especially when the tool is supposed to act in conjunction with a human user.

Second, there is little understanding of how AI interacts with humans within a healthcare context. For example, in a recent study[6], presenting results from a silent trial of an algorithm to predict obstructive hydronephrosis in children based on their renal ultrasounds, users exposed to the tool changed their clinical decision-making behavior based on their expectations of the model's outputs, effectively learning from the model. This is especially relevant with the advent of increasingly powerful large language models such as ChatGPT and foundation models, the behavior of which is comparatively less predictable or interpretable, in general.

Third, the evaluation of AI tools and devices should not be solely driven by operational measurements, such as whether the tool increases the productivity of clinicians or of the health system overall. Although these are important outcomes in the general quest for a more sustainable healthcare system, the evaluation of any model must account also actual benefit and potential harms for the individual or population at the other user's end.

Lastly, there is the risk of AI further increasing or creating new health disparities. The deployment of the most advanced tools depends on a digital infrastructure system that is simply not present in most countries. As new studies are designed and conducted, it will be important to consider the feasibility of implementation of an AI intervention where it will be most needed, including resource-limited settings. For instance, the widespread use of smartphones in low-income countries make app-based digital interventions a relatively easy way to bring distributed health assistance and support even in remote areas. Studies have already shown the potential of these applications for remote support for self-administered abortion[7] and antibiotic stewardship[8].

Many of these points are now reflected in the Responsible AI for Social and Ethical Healthcare (RAISE) statement[9], a consensus-based effort, organized by the Department of Biomedical Informatics at Harvard Medical School, involving many key stakeholders in the transition to AI-powered health care. We also want to see these principles more frequently reflected in the design of the studies that we publish.

Along those lines, we encourage submission of new research providing strong evidence to support the implementation of AI in healthcare, particularly in resources-limited settings, be it through clinical trials, prospective observational studies or real-world implementation and cost-effective research. We invite our authors to continue to partner with us to champion research that addresses these gaps and put AI forward as a healthcare democratization tool.

Editorial

References

1. World Health Organization. *Aging and Health*; https://go.nature.com/3v0P5P8 (2022).
2. Ng, A. Y. et al. *Nat. Med.* **29**, 3044–3049 (2023).
3. Lång, K. et al. *Lancet Oncol.* **24**, 936–944 (2023).
4. Dembrower, K. et al. *Lancet Digital Health* **5**, E703–E711 (2023).
5. Wynants, L. et al. *BMJ* **369**, m1328 (2020).
6. Kwong, J. C. C. et al. *New Engl. J. Med.* https://doi.org/10.1056/AIcs2300004 (2024).
7. Endler, M. et al. *Lancet* **400**, 670–679 (2022).
8. Tan, R. et al. *Nat. Med.* **30**, 76–84 (2024).
9. Goldberg, C. B. et al. *Nat. Med.* https://doi.org/10.1038/s41591-024-02853-7 (2024).