How will AI change health care delivery?

Professor Jim Warren
PhD, FACHI
Professor of Health Informatics
School of Computer Science
University of Auckland

Perhaps we haven't yet reached the point where computer-based Artificial Intelligence (AI) has overtaken humanity as the masters of the world, but there's a groundswell of sentiment that AI can now exceed human performance for almost any specialised task. In 2015, AI displaced a bastion of human mastery when a computer programme beat a human champion in Go, a game that had long held out against attempts to exceed the best human players, in part due to the size of the board and many possible moves. The winning computer program employed techniques that are hallmarks of the new wave of AI: using a deep neural network (a system of nodes and weighted connections with multiple layers between the inputs and the outputs); 'big data' (a comprehensive collection of transcripts of high-level human Go games in this case), and massive computation (notably, to learn from the data to recognise good moves and the value of board positions, reinforced by the equivalent of lifetimes of simulated games against itself.[1] AI has been an active field of research ever since digital electronic computers emerged after World War II, and while notoriously difficult to define, intertwined with concepts of rational and human-like thought and action AI can be taken simply as the attempt to build intelligent entities.[2] There's a tendency to move the threshold for what constitutes 'real AI' forward to exclude established innovations. For instance, AI accomplishments of past decades, such as the automated interpretation of electrocardiograms (ECGs), may now be seen merely as useful technology without much regard to the human-like nature of the task being accomplished.[3] But this new wave of AI based on deep learning has re-ignited both the imagination of the public in general and the health-care community, in particular in terms of the potential of AI to change our lives.

Learning to beat humans at games is an important part of AI research, not just as a publicity stunt, but for the insights learned in finding ways to out think humans at tasks that have attracted individuals to dedicate a lifetime of training to becoming experts (such as Grandmasters in chess). But AI has always had its applied side as well, including learning to imitate (or exceed) the performance of medical experts. Hard on the heels of the breakthrough in Go, a deep learning system was demonstrated to provide formidable sensitivity and specificity for detecting diabetic retinopathy in fundus images as compared to a panel of United States licensed ophthalmologists and ophthalmology senior residents.[4] The authors themselves took some care to point out limitations – the algorithm would not necessarily detect non-diabetic retinopathy lesions that were outside of its training data, nor would it be a replacement for a comprehensive eye exam – yet there is a temptation for the findings of this frequently-cited article (893 times on Google Scholar at 6 March 2019) to be consolidated simply as – with deep learning, AIs can now match specialists. A recent Journal of the American Medical Association editorial indicated that the new wave of AI is one of a series of technology-based advances, and makes a comparison to how computed tomography has become part of the radiology toolkit.[5] Nonetheless, the concluding words, 'artificial intelligence and deep learning are entering the mainstream of clinical medicine' and, 'physicians need to actively engage to adapt their practice', set a tone that we have reached a tipping point for AI in medical decision making. A medical student could be forgiven for feeling some anxiety, wondering just what a future with AI making better decisions than specialists implies for their role and the doctor-patient relationship, or how they might be expected to engage this phenomenon.

We can expect that AI systems for health application will continue to grow in diversity and effectiveness. 'Super-computing' is now readily available: the graphics processing units in the video cards of our home computers turn out to be superb number-crunchers for neural network algorithms; or we can rent scalable computing power through cloud computing services by Amazon, Google or others. Moreover, the ever-increasing permeation of health-care systems with computing has as its natural by-product a growing archive of electronic medical records ripe for analysis. While this AI boom is indeed likely to be transformative to health-care delivery, there are reasons to take the view that this change will be incremental, manageable, and (hopefully) on balance positive.

First, AI algorithms from deep learning are not so unlike computing capabilities that we have been using routinely in New Zealand for years. For example, PREDICT is simultaneously decision-support software and an ongoing, prospectively designed, open cohort study.[6] The PREDICT software integrates with the practice management system to retrieve patient data, with any remaining required data entered interactively to provide an individualised estimate of the probability of a cardiovascular disease (CVD) event in the next five years, along with treatment recommendations. Participant risk factors captured by software that is regularly linked to national databases included hospitalisations and deaths related to CVD, supporting ongoing research to improve the risk prediction – most recently, based on over 400,000 patient encounters in New Zealand from 2002–2015.[7] At the heart of the risk prediction is a regression model (specifically a Cox proportional-hazards model) that gives a particular weight to each risk factor. The model is structurally much simpler than a deep learning model, but has the advantage that the reasoning behind the model's recommendation is easily explained. Adding explanation ability to deep neural networks is an active research area.[8]

The experience, for patients and health-care professionals, in using a deep learning AI (at least one that has been appropriately developed and carefully tested) will be little different to that with PREDICT, which has integrated smoothly with the existing health-care system and professional roles.

Second, while AI will challenge the doctor-patient dynamic, information technology (IT) challenging the doctor-patient dynamic is nothing new. For over 25 years, the World Wide Web (the Web) has been democratising access to information. Patients are at liberty to bring into their consults printouts (or perhaps nowadays more likely to brandish their cell phone or tablet) with the latest research findings, as well as potentially questionable content biased by revenue generation motives. As the Web has become more sophisticated and IT reaches ever more intimately into our lives, so the diversity of ways patients may bring IT into their health care has grown, now including mobile apps, fitness trackers, and blog posts. An interesting example is PatientsLikeMe, a Web-based network where patients connect to others with the same disease and share experiences. Sharing of quantitative data is encouraged along with the organisation of research studies, for example to test the effectiveness of off-label uses of drugs.[9] In his book *The Patient Will See You Now*, Eric Topol describes medicine as having reached a 'Gutenberg moment', where new freedom of information is enabling health consumers to take a revolutionary degree of control of their health care.[10] Topol cites numerous Web and IT-mediated trends, including sharing of big data and direct-to-patient genetic test results (as exemplified by 23andme).[11] Meanwhile, mobile text-based services are slipping into the mainstream of evidence-based medicine. For instance, a program including motivational messages and behaviour-change techniques was shown to significantly improve smoking cessation rates at six months.[12] The package of intervention techniques and dialog strategies operationalised in this service in fact makes it a form of AI – one that can be recommended to a patient by a doctor, or that a consumer can find and download for themselves over the Web.

Third, health-care professionals can engage with, and encourage or moderate, the advance of AI by routinely asking questions of provenance. You may encounter AI-based decision support presented by a patient, or integrated with the systems you use in your Primary Health Organisation or District Health Board. In any event, you can query where it comes from – who is endorsing and distributing it, and what is their motivation (i.e, is it purely for profit through proliferation – licensing fees or banner-ad revenue – or is it publicly funded; is it endorsed by a medical body?). Is it part of the new wave of AI based on machine learning from big data? Or perhaps (as with the above smoking cessation example) the capability is a product of 'knowledge engineering', where techniques based on human experts have been deliberately selected. If it is based on data, then data from where and when? Does that data seem likely to be a good representation of your own patient population, or would there be obvious gaps (e.g. lacking Māori and Pacific cases)? Can the system be retrained on local data? Can the system offer explanations for its recommendations, or is it just a 'black box' that offers no specific insight for its assessment? Is there evidence of the system's effectiveness? If so, how has its performance been evaluated: in what context, on what population, over what duration, and particularly what was its performance compared to? If the answers to these questions are hard to find, you should be suspicious (or at least cautious); if the answers are unsatisfactory, you should actively communicate about the system's limitations.

To take the concept of engagement further, it is worth noting that Health Informatics is an established interdisciplinary field and a growing profession – this is the field that deals with methods of information processing and management in health care, including AI in health-care delivery. Membership in Health Informatics New Zealand (HINZ) is open to anyone with an interest in the field; HINZ events, particularly the annual national conference, are a great way to learn more about the field and meet the Health Informatics community. Several New Zealand universities offer postgraduate degrees in Health Informatics, and there are numerous options to study online with overseas universities (HINZ maintains a list of domestic and overseas study options: https://www.hinz.org.nz/page/EducationOptions). One can apply to become a member or fellow of the Australasian College of Health Informatics based on contribution to the field, and there now exists a training pathway to fellowship (https://www.achi.org.au/achi-fellowship-program/). While in this article I have taken a particularly medical/doctor centred view of the impact of AI on health-care delivery (given the nature of the journal), it is important to understand that the field is concerned with the whole health-care team; notably, nurses have been especially active in Health Informatics throughout its history. AI will influence and expand the capabilities of every type of professional associated with health care, as well as the health consumer.

The growing application of AI will add new and diverse inputs into the clinical context, but it will be just one more source of information to be considered in medical decision making. If you approach decision making as a shared process in partnership with patients, then they will be less likely to use Google to replace you!

References

1. Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, et al. Mastering the game of Go with deep neural networks and tree search. Nature. 2016;529(7587):484–9.

2. Russell S, Norvig P. Artificial intelligence: a modern approach. 3rd ed. Prentice-Hall; 2010. Chapter 1, Introduction; p. 1–33.

3. Yu KH, Kohane IS. Framing the challenges of artificial intelligence in medicine. BMJ Qual Saf. 2019;28(3):238–41.

4. Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA. 2016;316(22):2402–10.

5. Stead WW. Clinical implications and challenges of artificial intelligence and deep learning. JAMA. 2018;320(11):1107–8.

6. Wells S, Riddell T, Kerr A, Pylypchuk R, Chelimo C, Marshall R, et al. Cohort profile: the PREDICT cardiovascular disease cohort in New Zealand primary care (PREDICT-CVD 19). Int J Epidemiol. 2017;46(1):22.

7. Pylypchuk R, Wells S, Kerr A, Poppe K, Riddell T, Harwood M, et al. Cardiovascular disease risk prediction equations in 400 000 primary care patients in New Zealand: a derivation and validation study. Lancet. 2018;391(10133):1897–907.

8. Montavon G, Samek W, Müller K-R. Methods for interpreting and understanding deep neural networks. Digit Signal Process. 2018;73:1–15.

9. Frost J, Okun S, Vaughan T, Heywood J, Wicks P. Patient-reported outcomes as a source of evidence in off-label prescribing: analysis of data from PatientsLikeMe. J Med Internet Res 2011;13(1):e6.

10. Topol E. The patient will see you now: the future of medicine is in your hands. New York: Basic Books; 2015.

11. Stoekle HC, Mamzer-Bruneel MF, Vogt G, Herve C. 23andMe: a new two-sided data-banking market model. BMC Med Ethics. 2016;17:19.

12. Free C, Knight R, Robertson S, Whittaker R, Edwards P, Zhou W, et al. Smoking cessation support delivered via mobile phone text messaging (txt2stop): a single-blind, randomised trial. Lancet. 2011;378(9785):49–55.