

DIGITAL TECHNOLOGY The potential for artificial intelligence in healthcare

Authors: Thomas Davenport^A and Ravi Kalakota^B

ABSTRACT

The complexity and rise of data in healthcare means that artificial intelligence (AI) will increasingly be applied within the field. Several types of AI are already being employed by payers and providers of care, and life sciences companies. The key categories of applications involve diagnosis and treatment recommendations, patient engagement and adherence, and administrative activities. Although there are many instances in which AI can perform healthcare tasks as well or better than humans, implementation factors will prevent large-scale automation of healthcare professional jobs for a considerable period. Ethical issues in the application of AI to healthcare are also discussed.

KEYWORDS: Artificial intelligence, clinical decision support, electronic health record systems

Introduction

Artificial intelligence (AI) and related technologies are increasingly prevalent in business and society, and are beginning to be applied to healthcare. These technologies have the potential to transform many aspects of patient care, as well as administrative processes within provider, payer and pharmaceutical organisations.

There are already a number of research studies suggesting that AI can perform as well as or better than humans at key healthcare tasks, such as diagnosing disease. Today, algorithms are already outperforming radiologists at spotting malignant tumours, and guiding researchers in how to construct cohorts for costly clinical trials. However, for a variety of reasons, we believe that it will be many years before AI replaces humans for broad medical process domains. In this article, we describe both the potential that AI offers to automate aspects of care and some of the barriers to rapid implementation of AI in healthcare.

Types of AI of relevance to healthcare

Artificial intelligence is not one technology, but rather a collection of them. Most of these technologies have immediate relevance to the healthcare field, but the specific processes and tasks they

support vary widely. Some particular AI technologies of high importance to healthcare are defined and described below.

Machine learning – neural networks and deep learning

Machine learning is a statistical technique for fitting models to data and to 'learn' by training models with data. Machine learning is one of the most common forms of AI; in a 2018 Deloitte survey of 1,100 US managers whose organisations were already pursuing AI, 63% of companies surveyed were employing machine learning in their businesses.¹ It is a broad technique at the core of many approaches to AI and there are many versions of it.

In healthcare, the most common application of traditional machine learning is precision medicine – predicting what treatment protocols are likely to succeed on a patient based on various patient attributes and the treatment context.² The great majority of machine learning and precision medicine applications require a training dataset for which the outcome variable (eg onset of disease) is known; this is called supervised learning.

A more complex form of machine learning is the *neural network* – a technology that has been available since the 1960s has been well established in healthcare research for several decades³ and has been used for categorisation applications like determining whether a patient will acquire a particular disease. It views problems in terms of inputs, outputs and weights of variables or 'features' that associate inputs with outputs. It has been likened to the way that neurons process signals, but the analogy to the brain's function is relatively weak.

The most complex forms of machine learning involve *deep learning*, or neural network models with many levels of features or variables that predict outcomes. There may be thousands of hidden features in such models, which are uncovered by the faster processing of today's graphics processing units and cloud architectures. A common application of deep learning in healthcare is recognition of potentially cancerous lesions in radiology images.⁴ Deep learning is increasingly being applied to radiomics, or the detection of clinically relevant features in imaging data beyond what can be perceived by the human eye.⁵ Both radiomics and deep learning are most commonly found in oncology-oriented image analysis. Their combination appears to promise greater accuracy in diagnosis than the previous generation of automated tools for image analysis, known as computer-aided detection or CAD.

Deep learning is also increasingly used for speech recognition and, as such, is a form of natural language processing (NLP),

Authors: ^Apresident's distinguished professor of information technology and management, Babson College, Wellesley, USA; ^Bmanaging director, Deloitte Consulting, New York, USA

described below. Unlike earlier forms of statistical analysis, each feature in a deep learning model typically has little meaning to a human observer. As a result, the explanation of the model's outcomes may be very difficult or impossible to interpret.

Natural language processing

Making sense of human language has been a goal of AI researchers since the 1950s. This field, NLP, includes applications such as speech recognition, text analysis, translation and other goals related to language. There are two basic approaches to it: statistical and semantic NLP. Statistical NLP is based on machine learning (deep learning neural networks in particular) and has contributed to a recent increase in accuracy of recognition. It requires a large 'corpus' or body of language from which to learn.

In healthcare, the dominant applications of NLP involve the creation, understanding and classification of clinical documentation and published research. NLP systems can analyse unstructured clinical notes on patients, prepare reports (eg on radiology examinations), transcribe patient interactions and conduct conversational AI.

Rule-based expert systems

Expert systems based on collections of 'if-then' rules were the dominant technology for AI in the 1980s and were widely used commercially in that and later periods. In healthcare, they were widely employed for 'clinical decision support' purposes over the last couple of decades⁵ and are still in wide use today. Many electronic health record (EHR) providers furnish a set of rules with their systems today.

Expert systems require human experts and knowledge engineers to construct a series of rules in a particular knowledge domain. They work well up to a point and are easy to understand. However, when the number of rules is large (usually over several thousand) and the rules begin to conflict with each other, they tend to break down. Moreover, if the knowledge domain changes, changing the rules can be difficult and time-consuming. They are slowly being replaced in healthcare by more approaches based on data and machine learning algorithms.

Physical robots

Physical robots are well known by this point, given that more than 200,000 industrial robots are installed each year around the world. They perform pre-defined tasks like lifting, repositioning, welding or assembling objects in places like factories and warehouses, and delivering supplies in hospitals. More recently, robots have become more collaborative with humans and are more easily trained by moving them through a desired task. They are also becoming more intelligent, as other AI capabilities are being embedded in their 'brains' (really their operating systems). Over time, it seems likely that the same improvements in intelligence that we've seen in other areas of AI would be incorporated into physical robots.

Surgical robots, initially approved in the USA in 2000, provide 'superpowers' to surgeons, improving their ability to see, create precise and minimally invasive incisions, stitch wounds and so forth.⁶ Important decisions are still made by human surgeons, however. Common surgical procedures using robotic surgery include gynaecologic surgery, prostate surgery and head and neck surgery.

Robotic process automation

This technology performs structured digital tasks for administrative purposes, ie those involving information systems, as if they were a human user following a script or rules. Compared to other forms of AI they are inexpensive, easy to program and transparent in their actions. Robotic process automation (RPA) doesn't really involve robots – only computer programs on servers. It relies on a combination of workflow, business rules and 'presentation layer' integration with information systems to act like a semi-intelligent user of the systems. In healthcare, they are used for repetitive tasks like prior authorisation, updating patient records or billing. When combined with other technologies like image recognition, they can be used to extract data from, for example, faxed images in order to input it into transactional systems.⁷

We've described these technologies as individual ones, but increasingly they are being combined and integrated; robots are getting AI-based 'brains', image recognition is being integrated with RPA. Perhaps in the future these technologies will be so intermingled that composite solutions will be more likely or feasible.

Diagnosis and treatment applications

Diagnosis and treatment of disease has been a focus of AI since at least the 1970s, when MYCIN was developed at Stanford for diagnosing blood-borne bacterial infections.⁸ This and other early rule-based systems showed promise for accurately diagnosing and treating disease, but were not adopted for clinical practice. They were not substantially better than human diagnosticians, and they were poorly integrated with clinician workflows and medical record systems.

More recently, IBM's Watson has received considerable attention in the media for its focus on precision medicine, particularly cancer diagnosis and treatment. Watson employs a combination of machine learning and NLP capabilities. However, early enthusiasm for this application of the technology has faded as customers realised the difficulty of teaching Watson how to address particular types of cancer⁹ and of integrating Watson into care processes and systems.¹⁰ Watson is not a single product but a set of 'cognitive services' provided through application programming interfaces (APIs), including speech and language, vision, and machine learning-based data-analysis programs. Most observers feel that the Watson APIs are technically capable, but taking on cancer treatment was an overly ambitious objective. Watson and other proprietary programs have also suffered from competition with free 'open source' programs provided by some vendors, such as Google's TensorFlow.

Implementation issues with AI bedevil many healthcare organisations. Although rule-based systems incorporated within EHR systems are widely used, including at the NHS,¹¹ they lack the precision of more algorithmic systems based on machine learning. These rule-based clinical decision support systems are difficult to maintain as medical knowledge changes and are often not able to handle the explosion of data and knowledge based on genomic, proteomic, metabolic and other 'omic-based' approaches to care.

This situation is beginning to change, but it is mostly present in research labs and in tech firms, rather than in clinical practice. Scarcely a week goes by without a research lab claiming that it has developed an approach to using AI or big data to diagnose

and treat a disease with equal or greater accuracy than human clinicians. Many of these findings are based on radiological image analysis,¹² though some involve other types of images such as retinal scanning¹³ or genomic-based precision medicine.¹⁴ Since these types of findings are based on statistically-based machine learning models, they are ushering in an era of evidence- and probability-based medicine, which is generally regarded as positive but brings with it many challenges in medical ethics and patient/clinician relationships.¹⁵

Tech firms and startups are also working assiduously on the same issues. Google, for example, is collaborating with health delivery networks to build prediction models from big data to warn clinicians of high-risk conditions, such as sepsis and heart failure.¹⁶ Google, Enlitic and a variety of other startups are developing AI-derived image interpretation algorithms. Jvion offers a 'clinical success machine' that identifies the patients most at risk as well as those most likely to respond to treatment protocols. Each of these could provide decision support to clinicians seeking to find the best diagnosis and treatment for patients.

There are also several firms that focus specifically on diagnosis and treatment recommendations for certain cancers based on their genetic profiles. Since many cancers have a genetic basis, human clinicians have found it increasingly complex to understand all genetic variants of cancer and their response to new drugs and protocols. Firms like Foundation Medicine and Flatiron Health, both now owned by Roche, specialise in this approach.

Both providers and payers for care are also using 'population health' machine learning models to predict populations at risk of particular diseases¹⁷ or accidents¹⁸ or to predict hospital readmission.¹⁹ These models can be effective at prediction, although they sometimes lack all the relevant data that might add predictive capability, such as patient socio-economic status.

But whether rules-based or algorithmic in nature, AI-based diagnosis and treatment recommendations are sometimes challenging to embed in clinical workflows and EHR systems. Such integration issues have probably been a greater barrier to broad implementation of AI than any inability to provide accurate and effective recommendations; and many AI-based capabilities for diagnosis and treatment from tech firms are standalone in nature or address only a single aspect of care. Some EHR vendors have begun to embed limited AI functions (beyond rule-based clinical decision support) into their offerings,²⁰ but these are in the early stages. Providers will either have to undertake substantial integration projects themselves or wait until EHR vendors add more AI capabilities.

Patient engagement and adherence applications

Patient engagement and adherence has long been seen as the 'last mile' problem of healthcare – the final barrier between ineffective and good health outcomes. The more patients proactively participate in their own well-being and care, the better the outcomes – utilisation, financial outcomes and member experience. These factors are increasingly being addressed by big data and AI.

Providers and hospitals often use their clinical expertise to develop a plan of care that they know will improve a chronic or acute patient's health. However, that often doesn't matter if the patient fails to make the behavioural adjustment necessary, eg losing weight, scheduling a follow-up visit, filling prescriptions or complying with a treatment plan. Noncompliance – when

a patient does not follow a course of treatment or take the prescribed drugs as recommended – is a major problem.

In a survey of more than 300 clinical leaders and healthcare executives, more than 70% of the respondents reported having less than 50% of their patients highly engaged and 42% of respondents said less than 25% of their patients were highly engaged.²¹

If deeper involvement by patients results in better health outcomes, can AI-based capabilities be effective in personalising and contextualising care? There is growing emphasis on using machine learning and business rules engines to drive nuanced interventions along the care continuum.²² Messaging alerts and relevant, targeted content that provoke actions at moments that matter is a promising field in research.

Another growing focus in healthcare is on effectively designing the 'choice architecture' to nudge patient behaviour in a more anticipatory way based on real-world evidence. Through information provided by provider EHR systems, biosensors, watches, smartphones, conversational interfaces and other instrumentation, software can tailor recommendations by comparing patient data to other effective treatment pathways for similar cohorts. The recommendations can be provided to providers, patients, nurses, call-centre agents or care delivery coordinators.

Administrative applications

There are also a great many administrative applications in healthcare. The use of AI is somewhat less potentially revolutionary in this domain as compared to patient care, but it can provide substantial efficiencies. These are needed in healthcare because, for example, the average US nurse spends 25% of work time on regulatory and administrative activities.²³ The technology that is most likely to be relevant to this objective is RPA. It can be used for a variety of applications in healthcare, including claims processing, clinical documentation, revenue cycle management and medical records management.²⁴

Some healthcare organisations have also experimented with chatbots for patient interaction, mental health and wellness, and telehealth. These NLP-based applications may be useful for simple transactions like refilling prescriptions or making appointments. However, in a survey of 500 US users of the top five chatbots used in healthcare, patients expressed concern about revealing confidential information, discussing complex health conditions and poor usability.²⁵

Another AI technology with relevance to claims and payment administration is machine learning, which can be used for probabilistic matching of data across different databases. Insurers have a duty to verify whether the millions of claims are correct. Reliably identifying, analysing and correcting coding issues and incorrect claims saves all stakeholders – health insurers, governments and providers alike – a great deal of time, money and effort. Incorrect claims that slip through the cracks constitute significant financial potential waiting to be unlocked through data-matching and claims audits.

Implications for the healthcare workforce

There has been considerable attention to the concern that AI will lead to automation of jobs and substantial displacement of the workforce. A Deloitte collaboration with the Oxford Martin

Institute²⁶ suggested that 35% of UK jobs could be automated out of existence by AI over the next 10 to 20 years. Other studies have suggested that while some automation of jobs is possible, a variety of external factors other than technology could limit job loss, including the cost of automation technologies, labour market growth and cost, benefits of automation beyond simple labour substitution, and regulatory and social acceptance.²⁷ These factors might restrict actual job loss to 5% or less.

To our knowledge thus far there have been no jobs eliminated by AI in health care. The limited incursion of AI into the industry thus far, and the difficulty of integrating AI into clinical workflows and EHR systems, have been somewhat responsible for the lack of job impact. It seems likely that the healthcare jobs most likely to be automated would be those that involve dealing with digital information, radiology and pathology for example, rather than those with direct patient contact.²⁸

But even in jobs like radiologist and pathologist, the penetration of AI into these fields is likely to be slow. Even though, as we have argued, technologies like deep learning are making inroads into the capability to diagnose and categorise images, there are several reasons why radiology jobs, for example, will not disappear soon.²⁹

First, radiologists do more than read and interpret images. Like other AI systems, radiology AI systems perform single tasks. Deep learning models in labs and startups are trained for specific image recognition tasks (such as nodule detection on chest computed tomography or hemorrhage on brain magnetic resonance imaging). However, thousands of such narrow detection tasks are necessary to fully identify all potential findings in medical images, and only a few of these can be done by AI today. Radiologists also consult with other physicians on diagnosis and treatment, treat diseases (for example providing local ablative therapies) and perform image-guided medical interventions such as cancer biopsies and vascular stents (interventional radiology), define the technical parameters of imaging examinations to be performed (tailored to the patient's condition), relate findings from images to other medical records and test results, discuss procedures and results with patients, and many other activities.

Second, clinical processes for employing AI-based image work are a long way from being ready for daily use. Different imaging technology vendors and deep learning algorithms have different foci: the probability of a lesion, the probability of cancer, a nodule's feature or its location. These distinct foci would make it very difficult to embed deep learning systems into current clinical practice.

Third, deep learning algorithms for image recognition require 'labelled data' – millions of images from patients who have received a definitive diagnosis of cancer, a broken bone or other pathology. However, there is no aggregated repository of radiology images, labelled or otherwise.

Finally, substantial changes will be required in medical regulation and health insurance for automated image analysis to take off.

Similar factors are present for pathology and other digitally-oriented aspects of medicine. Because of them, we are unlikely to see substantial change in healthcare employment due to AI over the next 20 years or so. There is also the possibility that new jobs will be created to work with and to develop AI technologies. But static or increasing human employment also mean, of course, that AI technologies are not likely to substantially reduce the costs of medical diagnosis and treatment over that timeframe.

Ethical implications

Finally, there are also a variety of ethical implications around the use of AI in healthcare. Healthcare decisions have been made almost exclusively by humans in the past, and the use of smart machines to make or assist with them raises issues of accountability, transparency, permission and privacy.

Perhaps the most difficult issue to address given today's technologies is transparency. Many AI algorithms – particularly deep learning algorithms used for image analysis – are virtually impossible to interpret or explain. If a patient is informed that an image has led to a diagnosis of cancer, he or she will likely want to know why. Deep learning algorithms, and even physicians who are generally familiar with their operation, may be unable to provide an explanation.

Mistakes will undoubtedly be made by AI systems in patient diagnosis and treatment and it may be difficult to establish accountability for them. There are also likely to be incidents in which patients receive medical information from AI systems that they would prefer to receive from an empathetic clinician. Machine learning systems in healthcare may also be subject to algorithmic bias, perhaps predicting greater likelihood of disease on the basis of gender or race when those are not actually causal factors.³⁰

We are likely to encounter many ethical, medical, occupational and technological changes with AI in healthcare. It is important that healthcare institutions, as well as governmental and regulatory bodies, establish structures to monitor key issues, react in a responsible manner and establish governance mechanisms to limit negative implications. This is one of the more powerful and consequential technologies to impact human societies, so it will require continuous attention and thoughtful policy for many years.

The future of AI in healthcare

We believe that AI has an important role to play in the healthcare offerings of the future. In the form of machine learning, it is the primary capability behind the development of precision medicine, widely agreed to be a sorely needed advance in care. Although early efforts at providing diagnosis and treatment recommendations have proven challenging, we expect that AI will ultimately master that domain as well. Given the rapid advances in AI for imaging analysis, it seems likely that most radiology and pathology images will be examined at some point by a machine. Speech and text recognition are already employed for tasks like patient communication and capture of clinical notes, and their usage will increase.

The greatest challenge to AI in these healthcare domains is not whether the technologies will be capable enough to be useful, but rather ensuring their adoption in daily clinical practice. For widespread adoption to take place, AI systems must be approved by regulators, integrated with EHR systems, standardised to a sufficient degree that similar products work in a similar fashion, taught to clinicians, paid for by public or private payer organisations and updated over time in the field. These challenges will ultimately be overcome, but they will take much longer to do so than it will take for the technologies themselves to mature. As a result, we expect to see limited use of AI in clinical practice within 5 years and more extensive use within 10.

It also seems increasingly clear that AI systems will not replace human clinicians on a large scale, but rather will augment their efforts to care for patients. Over time, human clinicians may move

toward tasks and job designs that draw on uniquely human skills like empathy, persuasion and big-picture integration. Perhaps the only healthcare providers who will lose their jobs over time may be those who refuse to work alongside artificial intelligence. ■

References

- 1 Deloitte Insights. *State of AI in the enterprise*. Deloitte, 2018. www2.deloitte.com/content/dam/insights/us/articles/4780_State-of-AI-in-the-enterprise/AICognitiveSurvey2018_Infographic.pdf.
- 2 Lee SI, Celik S, Logsdon BA *et al*. A machine learning approach to integrate big data for precision medicine in acute myeloid leukemia. *Nat Commun* 2018;9:42.
- 3 Sordo M. *Introduction to neural networks in healthcare*. OpenClinical, 2002. www.openclinical.org/docs/int/neuralnetworks011.pdf
- 4 Fakoor R, Ladhak F, Nazi A, Huber M. *Using deep learning to enhance cancer diagnosis and classification*. A conference presentation. The 30th International Conference on Machine Learning, 2013.
- 5 Vial A, Stirling D, Field M *et al*. The role of deep learning and radiomic feature extraction in cancer-specific predictive modelling: a review. *Transl Cancer Res* 2018;7:803–16.
- 6 Davenport TH, Glaser J. Just-in-time delivery comes to knowledge management. *Harvard Business Review* 2002. <https://hbr.org/2002/07/just-in-time-delivery-comes-to-knowledge-management>.
- 7 Hussain A, Malik A, Halim MU, Ali AM. The use of robotics in surgery: a review. *Int J Clin Pract* 2014;68:1376–82.
- 8 Bush J. How AI is taking the scut work out of health care. *Harvard Business Review* 2018. <https://hbr.org/2018/03/how-ai-is-taking-the-scut-work-out-of-health-care>.
- 9 Buchanan BG, Shortliffe EH. *Rule-based expert systems: The MYCIN experiments of the Stanford heuristic programming project*. Reading: Addison Wesley, 1984.
- 10 Ross C, Swetlitz I. IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close. *Stat* 2017. www.statnews.com/2017/09/05/watson-ibm-cancer.
- 11 Davenport TH. *The AI Advantage*. Cambridge: MIT Press, 2018.
- 12 Right Care Shared Decision Making Programme, Capita. *Measuring shared decision making: A review of research evidence*. NHS, 2012. www.england.nhs.uk/wp-content/uploads/2013/08/7sdm-report.pdf.
- 13 Loria K. Putting the AI in radiology. *Radiology Today* 2018;19:10. www.radiologytoday.net/archive/rt0118p10.shtml.
- 14 Schmidt-Erfurth U, Bogunovic H, Sadeghipour A *et al*. Machine learning to analyze the prognostic value of current imaging biomarkers in neovascular age-related macular degeneration. *Ophthalmology Retina* 2018;2:24–30.
- 15 Aronson S, Rehm H. Building the foundation for genomic-based precision medicine. *Nature* 2015;526:336–42.
- 16 Rysavy M. Evidence-based medicine: A science of uncertainty and an art of probability. *Virtual Mentor* 2013;15:4–8.
- 17 Rajkomar A, Oren E, Chen K *et al*. Scalable and accurate deep learning with electronic health records. *npj Digital Medicine* 2018;1:18. www.nature.com/articles/s41746-018-0029-1.
- 18 Shimabukuro D, Barton CW, Feldman MD, Mataraso SJ, Das R. Effect of a machine learning-based severe sepsis prediction algorithm on patient survival and hospital length of stay: a randomised clinical trial. *BMJ Open Respir Res* 2017;4:e000234.
- 19 Aicha AN, Englebienne G, van Schooten KS, Pijnappels M, Kröse B. Deep learning to predict falls in older adults based on daily-Life trunk accelerometry. *Sensors* 2018;18:1654.
- 20 Low LL, Lee KH, Ong MEH *et al*. Predicting 30-Day readmissions: performance of the LACE index compared with a regression model among general medicine patients in Singapore. *Biomed Research International* 2015;2015:169870.
- 21 Davenport TH, Hongsermeier T, Mc Cord KA. Using AI to improve electronic health records. *Harvard Business Review* 2018. <https://hbr.org/2018/12/using-ai-to-improve-electronic-health-records>.
- 22 Volpp K, Mohta S. *Improved engagement leads to better outcomes, but better tools are needed*. *Insights Report*. NEJM Catalyst, 2016. <https://catalyst.nejm.org/patient-engagement-report-improved-engagement-leads-better-outcomes-better-tools-needed>.
- 23 Berg S. *Nudge theory explored to boost medication adherence*. Chicago: American Medical Association, 2018. www.ama-assn.org/delivering-care/patient-support-advocacy/nudge-theory-explored-boost-medication-adherence.
- 24 Commins J. *Nurses say distractions cut bedside time by 25%*. HealthLeaders, 2010. www.healthleadersmedia.com/nursing/nurses-say-distractions-cut-bedside-time-25.
- 25 Utermohlen K. *Four robotic process automation (RPA) applications in the healthcare industry*. Medium, 2018. <https://medium.com/@karl.uterhohlen/4-robotic-process-automation-rpa-applications-in-the-healthcare-industry-4d449b24b613>.
- 26 UserTesting. *Healthcare chatbot apps are on the rise but the overall customer experience (cx) falls short according to a UserTesting report*. San Francisco: UserTesting, 2019.
- 27 Deloitte. *From brawn to brains: The impact of technology on jobs in the UK*. Deloitte, 2015. www2.deloitte.com/content/dam/Deloitte/uk/Documents/Growth/deloitte-uk-insights-from-brawns-to-brain.pdf.
- 28 McKinsey Global Institute. *A future that works: automation, employment, and productivity*. McKinsey Global Institute, 2017. www.mckinsey.com/~/media/mckinsey/featured%20insights/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Executive-summary.ashx.
- 29 Davenport TH, Kirby J. *Only humans need apply: Winners and losers in the age of smart machines*. New York: HarperBusiness, 2016.
- 30 Davenport TH, Dreyer K. AI will change radiology, but it won't replace radiologists. *Harvard Business Review* 2018. <https://hbr.org/2018/03/ai-will-change-radiology-but-it-wont-replace-radiologists>.
- 31 Char DS, Shah NH, Magnus D. Implementing machine learning in health care – addressing ethical challenges. *N Engl J Med* 2018;378:981–3.

Address for correspondence: Prof Thomas Davenport, president's distinguished professor of information technology and management, Babson College, 231 Forest Street, Wellesley, MA 02457, USA. Email: tdavenport@babson.edu