

DIGITAL TECHNOLOGY Artificial intelligence in healthcare: transforming the practice of medicine

Authors: Junaid Bajwa,^A Usman Munir,^B Aditya Nori^C and Bryan Williams^D

ABSTRACT

Artificial intelligence (AI) is a powerful and disruptive area of computer science, with the potential to fundamentally transform the practice of medicine and the delivery of healthcare. In this review article, we outline recent breakthroughs in the application of AI in healthcare, describe a roadmap to building effective, reliable and safe AI systems, and discuss the possible future direction of AI augmented healthcare systems.

KEYWORDS: AI, digital health

DOI: 10.7861/fhj.2021-0095

Introduction

Healthcare systems around the world face significant challenges in achieving the ‘quadruple aim’ for healthcare: improve population health, improve the patient’s experience of care, enhance caregiver experience and reduce the rising cost of care.^{1–3} Ageing populations, growing burden of chronic diseases and rising costs of healthcare globally are challenging governments, payers, regulators and providers to innovate and transform models of healthcare delivery. Moreover, against a backdrop now catalysed by the global pandemic, healthcare systems find themselves challenged to ‘perform’ (deliver effective, high-quality care) and ‘transform’ care at scale by leveraging real-world data driven insights directly into patient care. The pandemic has also highlighted the shortages in healthcare workforce and inequities in the access to care, previously articulated by The King’s Fund and the World Health Organization (Box 1).^{4,5}

The application of technology and artificial intelligence (AI) in healthcare has the potential to address some of these supply-and-demand challenges. The increasing availability of multi-modal data (genomics, economic, demographic, clinical and phenotypic) coupled with technology innovations in mobile, internet of things (IoT), computing power and data security herald

Authors: ^Achief medical scientist, Microsoft Research, Cambridge, UK; ^Bresearch program manager, Microsoft Research, Cambridge, UK; ^Chead of health intelligence, Microsoft Research, Cambridge, UK; ^Dchair of medicine, University College London, London, UK and director, NIHR UCLH Biomedical Research Centre, London, UK

Box 1. Workforce challenges in the next decade

By 2030, the gap between supply of and demand for staff employed by NHS trusts could increase to almost 250,000 full-time equivalent posts.⁴

Based on the current trends and needs of the global population by 2030, the world will have 18 million fewer healthcare professionals (especially marked differences in the developing world), including 5 million fewer doctors than society will require.⁵

a moment of convergence between healthcare and technology to fundamentally transform models of healthcare delivery through AI-augmented healthcare systems.

In particular, cloud computing is enabling the transition of effective and safe AI systems into mainstream healthcare delivery. Cloud computing is providing the computing capacity for the analysis of considerably large amounts of data, at higher speeds and lower costs compared with historic ‘on premises’ infrastructure of healthcare organisations. Indeed, we observe that many technology providers are increasingly seeking to partner with healthcare organisations to drive AI-driven medical innovation enabled by cloud computing and technology-related transformation (Box 2).^{6–8}

Here, we summarise recent breakthroughs in the application of AI in healthcare, describe a roadmap to building effective AI systems and discuss the possible future direction of AI augmented healthcare systems.


Box 2. Quotes from technology leaders

Satya Nadella, chief executive officer, Microsoft: ‘AI is perhaps the most transformational technology of our time, and healthcare is perhaps AI’s most pressing application.’⁶

Tim Cook, chief executive officer, Apple: ‘[Healthcare] is a business opportunity ... if you look at it, medical health activity is the largest or second-largest component of the economy.’⁷

Google Health: ‘We think that AI is poised to transform medicine, delivering new, assistive technologies that will empower doctors to better serve their patients. Machine learning has dozens of possible application areas, but healthcare stands out as a remarkable opportunity to benefit people.’⁸

Fig 1. Multi-step, iterative approach to build effective and reliable AI-augmented systems in healthcare.



What is artificial intelligence?

Simply put, AI refers to the science and engineering of making intelligent machines, through algorithms or a set of rules, which the machine follows to mimic human cognitive functions, such as learning and problem solving.⁹ AI systems have the potential to anticipate problems or deal with issues as they come up and, as such, operate in an intentional, intelligent and adaptive manner.¹⁰ AI's strength is in its ability to learn and recognise patterns and relationships from large multidimensional and multimodal datasets; for example, AI systems could translate a patient's entire medical record into a single number that represents a likely diagnosis.^{11,12} Moreover, AI systems are dynamic and autonomous, learning and adapting as more data become available.¹³

AI is not one ubiquitous, universal technology, rather, it represents several subfields (such as machine learning and deep learning) that, individually or in combination, add intelligence to applications. Machine learning (ML) refers to the study of algorithms that allow computer programs to automatically improve through experience.¹⁴ ML itself may be categorised as 'supervised', 'unsupervised' and 'reinforcement learning' (RL), and there is ongoing research in various sub-fields including 'semi-supervised', 'self-supervised' and 'multi-instance' ML.

- > Supervised learning leverages labelled data (annotated information); for example, using labelled X-ray images of known tumours to detect tumours in new images.¹⁵
- > 'Unsupervised learning' attempts to extract information from data without labels; for example, categorising groups of patients with similar symptoms to identify a common cause.¹⁶
- > In RL, computational agents learn by trial and error, or by expert demonstration. The algorithm learns by developing a strategy to maximise rewards. Of note, major breakthroughs in AI in recent years have been based on RL.
- > Deep learning (DL) is a class of algorithms that learns by using a large, many-layered collection of connected processes and exposing these processors to a vast set of examples. DL has emerged as the predominant method in AI today driving improvements in areas such as image and speech recognition.^{17,18}

How to build effective and trusted AI-augmented healthcare systems?

Despite more than a decade of significant focus, the use and adoption of AI in clinical practice remains limited, with many AI products for healthcare still at the design and develop stage.^{19–22} While there are different ways to build AI systems for healthcare, far too often there are attempts to force square pegs into round

holes ie find healthcare problems to apply AI solutions to without due consideration to local context (such as clinical workflows, user needs, trust, safety and ethical implications).

We hold the view that AI amplifies and augments, rather than replaces, human intelligence. Hence, when building AI systems in healthcare, it is key to not replace the important elements of the human interaction in medicine but to focus it, and improve the efficiency and effectiveness of that interaction. Moreover, AI innovations in healthcare will come through an in-depth, human-centred understanding of the complexity of patient journeys and care pathways.

In Fig 1, we describe a problem-driven, human-centred approach, adapted from frameworks by Wiens *et al*, Care and Sendak to building effective and reliable AI-augmented healthcare systems.^{23–25}

Design and develop

The first stage is to design and develop AI solutions for the right problems using a human-centred AI and experimentation approach and engaging appropriate stakeholders, especially the healthcare users themselves.

Stakeholder engagement and co-creation

Build a multidisciplinary team including computer and social scientists, operational and research leadership, and clinical stakeholders (physician, caregivers and patients) and subject experts (eg for biomedical scientists) that would include authorisers, motivators, financiers, conveners, connectors, implementers and champions.²⁶ A multi-stakeholder team brings the technical, strategic, operational expertise to define problems, goals, success metrics and intermediate milestones.

Human-centred AI

A human-centred AI approach combines an ethnographic understanding of health systems, with AI. Through user-designed research, first understand the key problems (we suggest using a qualitative study design to understand 'what is the problem', 'why is it a problem', 'to whom does it matter', 'why has it not been addressed before' and 'why is it not getting attention') including the needs, constraints and workflows in healthcare organisations, and the facilitators and barriers to the integration of AI within the clinical context. After defining key problems, the next step is to identify which problems are appropriate for AI to solve, whether there is availability of applicable datasets to build and later evaluate AI. By contextualising algorithms in an existing workflow, AI systems would operate within existing norms and practices

to ensure adoption, providing appropriate solutions to existing problems for the end user.

Experimentation

The focus should be on piloting of new stepwise experiments to build AI tools, using tight feedback loops from stakeholders to facilitate rapid experiential learning and incremental changes.²⁷ The experiments would allow the trying out of new ideas simultaneously, exploring to see which one works, learn what works and what doesn't, and why.²⁸ Experimentation and feedback will help to elucidate the purpose and intended uses for the AI system: the likely end users and the potential harm and ethical implications of AI system to them (for instance, data privacy, security, equity and safety).

Evaluate and validate

Next, we must iteratively evaluate and validate the predictions made by the AI tool to test how well it is functioning. This is critical, and evaluation is based on three dimensions: statistical validity, clinical utility and economic utility.

- > Statistical validity is understanding the performance of AI on metrics of accuracy, reliability, robustness, stability and calibration. High model performance on retrospective, *in silico* settings is not sufficient to demonstrate clinical utility or impact.
- > To determine clinical utility, evaluate the algorithm in a real-time environment on a hold-out and temporal validation set (eg longitudinal and external geographic datasets) to demonstrate clinical effectiveness and generalisability.²⁵
- > Economic utility quantifies the net benefit relative to the cost from the investment in the AI system.

Scale and diffuse

Many AI systems are initially designed to solve a problem at one healthcare system based on the patient population specific to that location and context. Scale up of AI systems requires special attention to deployment modalities, model updates, the regulatory system, variation between systems and reimbursement environment.

Monitor and maintain

Even after an AI system has been deployed clinically, it must be continually monitored and maintained to monitor for risks and adverse events using effective post-market surveillance. Healthcare organisations, regulatory bodies and AI developers should cooperate to collate and analyse the relevant datasets for AI performance, clinical and safety-related risks, and adverse events.²⁹

What are the current and future use cases of AI in healthcare?

AI can enable healthcare systems to achieve their 'quadruple aim' by democratising and standardising a future of connected and AI augmented care, precision diagnostics, precision therapeutics and, ultimately, precision medicine (Table 1).³⁰ Research in the application of AI healthcare continues to accelerate rapidly, with potential use cases being demonstrated across the healthcare sector (both physical and mental health) including drug discovery, virtual clinical consultation, disease diagnosis, prognosis, medication management and health monitoring.

We describe a non-exhaustive suite of AI applications in healthcare in the near term, medium term and longer term, for the

Table 1. Widescale adoption and application of artificial intelligence in healthcare

Timeline	Connected/ augmented care	Precision diagnostics	Precision therapeutics	Precision Medicine	Summary
Short term: 0–5 years	Internet of things in healthcare Virtual assistants Augmented telehealth Personalised mental health support	Precision imaging (eg diabetic retinopathy and radiotherapy planning)	CRISPR (increasing use)	Digital and AI enabled research hospitals ³⁰	AI automates time consuming, high-volume repetitive tasks, especially within precision imaging
Medium-term: 5–10 years	Ambient intelligence in healthcare	Large-scale adoption and scale-up of precision imaging	Synthetic biology Immunomics	Customisation of healthcare Robotic assisted therapies	AI uses multi-modal datasets to drive precision therapeutics
Long term: >10 years	Autonomous virtual health assistants, delivering predictive and anticipatory care Networked and connected care organisations (single digital infrastructure)	Holographic and hybrid imaging Holomics (integrated genomic/radiomic/proteomic/clinical/immunohistochemical data)	Genomics medicine AI driven drug discovery	New curative treatments AI empowered healthcare professionals (eg digital twins)	AI enables healthcare systems to achieve a state of precision medicine through AI-augmented healthcare and connected care

Timings are illustrative to widescale adoption of the proposed innovation taking into account challenges / regulatory environment / use at scale.

potential capabilities of AI to augment, automate and transform medicine.

AI today (and in the near future)

Currently, AI systems are not reasoning engines ie cannot reason the same way as human physicians, who can draw upon 'common sense' or 'clinical intuition and experience'.¹² Instead, AI resembles a signal translator, translating patterns from datasets. AI systems today are beginning to be adopted by healthcare organisations to automate time consuming, high volume repetitive tasks. Moreover, there is considerable progress in demonstrating the use of AI in precision diagnostics (eg diabetic retinopathy and radiotherapy planning).

AI in the medium term (the next 5–10 years)

In the medium term, we propose that there will be significant progress in the development of powerful algorithms that are efficient (eg require less data to train), able to use unlabelled data, and can combine disparate structured and unstructured data including imaging, electronic health data, multi-omic, behavioural and pharmacological data. In addition, healthcare organisations and medical practices will evolve from being *adopters* of AI platforms, to becoming *co-innovators* with technology partners in the development of novel AI systems for precision therapeutics.

AI in the long term (>10 years)

In the long term, AI systems will become more *intelligent*, enabling AI healthcare systems achieve a state of precision medicine through AI-augmented healthcare and connected care. Healthcare will shift from the traditional one-size-fits-all form of medicine to a preventative, personalised, data-driven disease management model that achieves improved patient outcomes (improved patient and clinical experiences of care) in a more cost-effective delivery system.

Connected/augmented care

AI could significantly reduce inefficiency in healthcare, improve patient flow and experience, and enhance caregiver experience and patient safety through the care pathway; for example, AI could be applied to the remote monitoring of patients (eg intelligent telehealth through wearables/sensors) to identify and provide timely care of patients at risk of deterioration.

In the long term, we expect that healthcare clinics, hospitals, social care services, patients and caregivers to be all connected to a single, interoperable digital infrastructure using passive sensors in combination with ambient intelligence.³¹ Following are two AI applications in connected care.

Virtual assistants and AI chatbots

AI chatbots (such as those used in Babylon (www.babylonhealth.com) and Ada (<https://ada.com>)) are being used by patients to identify symptoms and recommend further actions in community and primary care settings. AI chatbots can be integrated with wearable devices such as smartwatches to provide insights to both patients and caregivers in improving their behaviour, sleep and general wellness.

Ambient and intelligent care

We also note the emergence of ambient sensing without the need for any peripherals.

- > Emerald (www.emeraldinno.com): a wireless, touchless sensor and machine learning platform for remote monitoring of sleep, breathing and behaviour, founded by Massachusetts Institute of Technology faculty and researchers.
- > Google nest: claiming to monitor sleep (including sleep disturbances like cough) using motion and sound sensors.³²
- > A recently published article exploring the ability to use smart speakers to contactlessly monitor heart rhythms.³³
- > Automation and ambient clinical intelligence: AI systems leveraging natural language processing (NLP) technology have the potential to automate administrative tasks such as documenting patient visits in electronic health records, optimising clinical workflow and enabling clinicians to focus more time on caring for patients (eg Nuance Dragon Ambient eXperience (www.nuance.com/healthcare/ambient-clinical-intelligence.html)).

Precision diagnostics

Diagnostic imaging

The automated classification of medical images is the leading AI application today. A recent review of AI/ML-based medical devices approved in the USA and Europe from 2015–2020 found that more than half (129 (58%) devices in the USA and 126 (53%) devices in Europe) were approved or CE marked for radiological use.³⁴ Studies have demonstrated AI's ability to meet or exceed the performance of human experts in image-based diagnoses from several medical specialties including pneumonia in radiology (a convolutional neural network trained with labelled frontal chest X-ray images outperformed radiologists in detecting pneumonia), dermatology (a convolutional neural network was trained with clinical images and was found to classify skin lesions accurately), pathology (one study trained AI algorithms with whole-slide pathology images to detect lymph node metastases of breast cancer and compared the results with those of pathologists) and cardiology (a deep learning algorithm diagnosed heart attack with a performance comparable with that of cardiologists).^{35–38}

We recognise that there are some exemplars in this area in the NHS (eg University of Leeds Virtual Pathology Project and the National Pathology Imaging Co-operative) and expect widescale adoption and scaleup of AI-based diagnostic imaging in the medium term.³⁹ We provide two use cases of such technologies.

Diabetic retinopathy screening

Key to reducing preventable, diabetes-related vision loss worldwide is screening individuals for detection and the prompt treatment of diabetic retinopathy. However, screening is costly given the substantial number of diabetes patients and limited manpower for eye care worldwide.⁴⁰ Research studies on automated AI algorithms for diabetic retinopathy in the USA, Singapore, Thailand and India have demonstrated robust diagnostic performance and cost effectiveness.^{41–44} Moreover, Centers for Medicare & Medicaid Services approved Medicare reimbursement for the use of Food and Drug Administration approved AI algorithm 'IDx-DR', which demonstrated 87%




Fig 2. Potential applications for the InnerEye deep learning toolkit include quantitative radiology for monitoring tumour progression, planning for surgery and radiotherapy planning.⁴⁷

sensitivity and 90% specificity for detecting more-than-mild diabetic retinopathy.⁴⁵

Improving the precision and reducing waiting timings for radiotherapy planning

An important AI application is to assist clinicians for image preparation and planning tasks for radiotherapy cancer treatment. Currently, segmentation of the images is time consuming and laborious task, performed manually by an oncologist using specially designed software to draw contours around the regions of interest. The AI-based InnerEye open-source technology can cut this preparation time for head and neck, and prostate cancer by up to 90%, meaning that waiting times for starting potentially life-saving radiotherapy treatment can be dramatically reduced (Fig 2).^{46,47}

Precision therapeutics

To make progress towards precision therapeutics, we need to considerably improve our understanding of disease. Researchers globally are exploring the cellular and molecular basis of disease, collecting a range of multimodal datasets that can lead to digital and biological biomarkers for diagnosis, severity and progression. Two important future AI applications include immunomics / synthetic biology and drug discovery.

Immunomics and synthetic biology

Through the application of AI tools on multimodal datasets in the future, we may be able to better understand the cellular basis of disease and the clustering of diseases and patient populations to provide more targeted preventive strategies, for example, using immunomics to diagnose and better predict care and treatment options. This will be revolutionary for multiple standards of care, with particular impact in the cancer, neurological and rare disease space, personalising the experience of care for the individual.

AI-driven drug discovery

AI will drive significant improvement in clinical trial design and optimisation of drug manufacturing processes, and, in general, any combinatorial optimisation process in healthcare could be replaced by AI. We have already seen the beginnings of this with the recent announcements by DeepMind and AlphaFold, which now sets the stage for better understanding disease processes, predicting protein structures and developing more targeted therapeutics (for both rare and more common diseases; Fig 3).^{48,49}

Precision medicine

New curative therapies

Over the past decade, synthetic biology has produced developments like CRISPR gene editing and some personalised cancer therapies. However, the life cycle for developing such advanced therapies is still extremely inefficient and expensive.

In future, with better access to data (genomic, proteomic, glycomic, metabolomic and bioinformatic), AI will allow us to handle far more systematic complexity and, in turn, help us transform the way we understand, discover and affect biology. This will improve the efficiency of the drug discovery process by helping better predict early which agents are more likely to be effective and also better anticipate adverse drug effects, which have often thwarted the further development of otherwise effective drugs at a costly late stage in the development process. This, in turn will democratise access to novel advanced therapies at a lower cost.

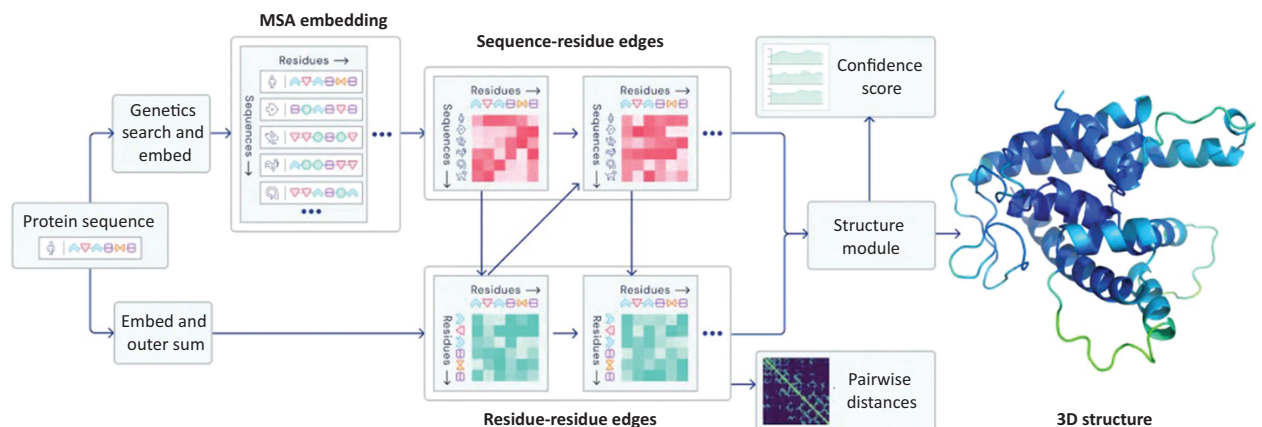


Fig 3. An overview of the main neural network model architecture for AlphaFold.⁴⁹ MSA = multiple sequence alignment.

AI empowered healthcare professionals

In the longer term, healthcare professionals will leverage AI in augmenting the care they provide, allowing them to provide safer, standardised and more effective care at the top of their licence; for example, clinicians could use an 'AI digital consult' to examine 'digital twin' models of their patients (a truly 'digital and biomedical' version of a patient), allowing them to 'test' the effectiveness, safety and experience of an intervention (such as a cancer drug) in the digital environment prior to delivering the intervention to the patient in the real world.

Challenges

We recognise that there are significant challenges related to the wider adoption and deployment of AI into healthcare systems. These challenges include, but are not limited to, data quality and access, technical infrastructure, organisational capacity, and ethical and responsible practices in addition to aspects related to safety and regulation. Some of these issues have been covered, but others go beyond the scope of this current article.

Conclusion and key recommendations

Advances in AI have the potential to transform many aspects of healthcare, enabling a future that is more personalised, precise, predictive and portable. It is unclear if we will see an incremental adoption of new technologies or radical adoption of these technological innovations, but the impact of such technologies and the digital renaissance they bring requires health systems to consider how best they will adapt to the changing landscape. For the NHS, the application of such technologies truly has the potential to release time for care back to healthcare professionals, enabling them to focus on what matters to their patients and, in the future, leveraging a globally democratised set of data assets comprising the 'highest levels of human knowledge' to 'work at the limits of science' to deliver a common high standard of care, wherever and whenever it is delivered, and by whoever.⁵⁰ Globally, AI could become a key tool for improving health equity around the world.

As much as the last 10 years have been about the roll out of digitisation of health records for the purposes of efficiency (and in some healthcare systems, billing/reimbursement), the next 10 years will be about the insight and value society can gain from these digital assets, and how these can be translated into driving better clinical outcomes with the assistance of AI, and the subsequent creation of novel data assets and tools. It is clear that we are at a turning point as it relates to the convergence of the practice of medicine and the application of technology, and although there are multiple opportunities, there are formidable challenges that need to be overcome as it relates to the real world and the scale of implementation of such innovation. A key to delivering this vision will be an expansion of translational research in the field of healthcare applications of artificial intelligence. Alongside this, we need investment into the upskilling of a healthcare workforce and future leaders that are digitally enabled, and to understand and embrace, rather than being intimidated by, the potential of an AI-augmented healthcare system.

Healthcare leaders should consider (as a minimum) these issues when planning to leverage AI for health:

- > processes for ethical and responsible access to data: healthcare data is highly sensitive, inconsistent, siloed and not optimised

for the purposes of machine learning development, evaluation, implementation and adoption

- > access to domain expertise / prior knowledge to make sense and create some of the rules which need to be applied to the datasets (to generate the necessary insight)
- > access to sufficient computing power to generate decisions in real time, which is being transformed exponentially with the advent of cloud computing
- > research into implementation: critically, we must consider, explore and research issues which arise when you take the algorithm and put it in the real world, building 'trusted' AI algorithms embedded into appropriate workflows. ■

References

- 1 Berwick DM, Nolan TW, Whittington J. The Triple Aim: Care, health, and cost. *Health Affairs* 2008;27:759–69.
- 2 Bodenheimer T, Sinsky C. From triple to quadruple aim: care of the patient requires care of the provider. *Ann Fam Med* 2014;12: 573–6.
- 3 Feeley D. *The Triple Aim or the Quadruple Aim? Four Points to Help Set Your Strategy*. Institute for Healthcare Improvement, 2017. www.ihl.org/communities/blogs/the-triple-aim-or-the-quadruple-aim-four-points-to-help-set-your-strategy.
- 4 The Health Foundation, Nuffield Trust, The King's Fund. *The health care workforce in England: make or break?* The King's Fund, 2018.
- 5 World Health Organization. *Working for health and growth: Investing in the health workforce*. WHO, 2016. <http://apps.who.int/iris/bitstream/10665/250047/1/9789241511308-eng.pdf> [Accessed 31 January 2020].
- 6 Satya Nadella announces strategic collaboration with Novartis. You Tube, 2019. www.youtube.com/watch?v=wMfsQE-D2q4
- 7 Lashinsky A. *Tim Cook on how Apple champions the environment, education, and health care*. Fortune, 2017.
- 8 Turea M. *How the 'Big 4' tech companies are leading healthcare innovation*. Healthcare Weekly, 2019.
- 9 McCarthy J. *What is artificial intelligence?* John McCarthy, 1998.
- 10 Shukla SS, Jaiswal V. Applicability of artificial intelligence in different fields of life. *IJSER* 2013;1:28–35.
- 11 Deng J, Dong W, Socher R et al. Imagenet: a large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition* 2009:248–55.
- 12 Quinn TP, Senadeera M, Jacobs S, Coghlan S, Le V. Trust and medical AI: the challenges we face and the expertise needed to overcome them. *J Am Med Inform Assoc* 2021;28:890–4.
- 13 Binns R, Gallo V. *Trade-offs*. Information Commissioner's Office, 2019. <https://ico.org.uk/about-the-ico/news-and-events/ai-blog-trade-offs>
- 14 Mitchell T. *Machine learning*. McGraw Hill, 1997. www.cs.cmu.edu/afs/cs.cmu.edu/user/mitchell/ftp/mlbook.html
- 15 Reardon S. Rise of robot radiologists. *Nature* 2019;576:S54–8.
- 16 Lasko TA, Denny JC, Levy MA. Computational phenotype discovery using unsupervised feature learning over noisy, sparse, and irregular clinical data. *PLoS One* 2013;8:e66341.
- 17 The Royal Society. *Machine learning: the power and promise of computers that learn by example*. The Royal Society, 2017.
- 18 LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 2015;521: 436–44.
- 19 Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med* 2019;25:44–56.
- 20 Kelly CJ, Karthikesalingam A, Suleyman M, Corrado G, King D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 2019;17:195.
- 21 Panch T, Mattie H, Celi LA. The 'inconvenient truth' about AI in healthcare. *NPJ Digit Med* 2019;2:77.
- 22 NHSX. *Artificial intelligence: How to get it right*. NHS, 2019.

- 23 Wiens J, Saria S, Sendak M *et al.* Do no harm: a roadmap for responsible machine learning for health care. *Nat Med* 2019;25:1337–40.
- 24 United States Government Accountability Office. *Artificial intelligence in health care: Benefits and challenges of technologies to augment patient care.* GAO, 2020.
- 25 Sendak MP, D'Arcy J, Kashyap S *et al.* A path for translation of machine learning products into healthcare delivery. *EMJ Innov* 2020;10:19-00172.
- 26 Andrews M, McConnell J, Wescott A. *Development as leadership-led change: working paper series RWP10-009.* Harvard University, 2010.
- 27 Andrews M, Pritchett L, Woolcock M. *Escaping capability traps through problem-driven iterative adaptation (PDIA): WIDER working paper 2012/064.* UNU-WIDER, 2012.
- 28 Andrews M. *Who really leads development? WIDER working paper 2013/092.* UNU-WIDER, 2013.
- 29 Davahli MR, Karwowski W, Fiok K, Wan T, Parsaei HR. Controlling safety of artificial intelligence-based systems in healthcare. *Symmetry* 2021;13:102.
- 30 Nachev P, Herron D, McNally N, Rees G, Williams B. Redefining the research hospital. *NPJ Digit Med* 2019;2:119.
- 31 Haque A, Milstein A, Fei-Fei L. Illuminating the dark spaces of healthcare with ambient intelligence. *Nature* 2020;585:193–202.
- 32 Muoio D. *Google's next-gen Nest Hub debuts with contactless sleep monitoring and analysis features.* Mobi Health News, 2021. www.mobihealthnews.com/news/googles-next-gen-nest-hub-debuts-contactless-sleep-monitoring-and-analysis-features
- 33 Wang A, Nguyen D, Sridhar AR *et al.* Using smart speakers to contactlessly monitor heart rhythms. *Commun Biol* 2021;4:319.
- 34 Muehlematter UJ, Daniore P, Vokinger KN. Approval of artificial intelligence and machine learning-based medical devices in the USA and Europe (2015–20): a comparative analysis. *Lancet Digital Health* 2021;3:e195–203.
- 35 Wang X, Peng Y, Lu L *et al.* Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. *IEEE CVPR* 2017:2097–106.
- 36 Esteva A, Robicquet A, Ramsundar B *et al.* A guide to deep learning in healthcare. *Nat Med* 2019;25:24–9.
- 37 Bejnordi BE, Veta M, Van Diest PJ *et al.* Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA* 2017;318:2199–210.
- 38 Strodthoff N, Strodthoff C. Detecting and interpreting myocardial infarction using fully convolutional neural networks. *Physiological Measurement* 2019;40:015001.
- 39 University of Leeds. *NPIC - Northern Pathology Imaging Co-operative.* University of Leeds, 2020. www.virtualpathology.leeds.ac.uk/npic
- 40 Belleme V, Lim ZW, Lim G *et al.* Artificial intelligence using deep learning to screen for referable and vision-threatening diabetic retinopathy in Africa: a clinical validation study. *Lancet Digit Health* 2019;1:e35–44.
- 41 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;316:2402–10.
- 42 Ting DSW, Pasquale LR, Peng L *et al.* Artificial intelligence and deep learning in ophthalmology. *Br J Ophthalmol* 2019;103:167–75.
- 43 Raumviboonsuk P, Krause J, Chotcomwongse P *et al.* Deep learning versus human graders for classifying diabetic retinopathy severity in a nationwide screening program. *NPJ Digit Med* 2019;2:25.
- 44 Xie Y, Nguyen QD, Hamzah H *et al.* Artificial intelligence for teleophthalmology-based diabetic retinopathy screening in a national programme: an economic analysis modelling study. *Lancet Digit Health* 2020;2:e240–9.
- 45 Simonite T. *The US government will pay doctors to use these AI algorithms.* Wired, 2020. www.wired.com/story/us-government-pay-doctors-use-ai-algorithms
- 46 Oktay O, Nanavati J, Schwaighofer A *et al.* Evaluation of deep learning to augment image-guided radiotherapy for head and neck and prostate cancers. *JAMA Netw Open* 2020;3:e2027426.
- 47 Alvarez-Valle J, Moore GJ. *Project InnerEye open-source deep learning toolkit: Democratizing medical imaging AI.* Microsoft, 2020. www.microsoft.com/en-us/research/blog/project-innereye-open-source-deep-learning-toolkit-democratizing-medical-imaging-ai
- 48 Senior AW, Evans R, Jumper J *et al.* Improved protein structure prediction using potentials from deep learning. *Nature* 2020;577:706–10.
- 49 The AlphaFold team. *AlphaFold: a solution to a 50-year-old grand challenge in biology.* DeepMind, 2020. <https://deepmind.com/blog/article/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology>
- 50 Department of Health and Social Care. *NHS Constitution or England.* DHSC, 2012 www.gov.uk/government/publications/the-nhs-constitution-for-england

**Address for correspondence: Dr Junaid Bajwa, Microsoft Research, 21 Station Road, Cambridge CB1 2FB, UK.
Email: junaid.bajwa@microsoft.com**