



OPEN ACCESS

# Medicine and the rise of the robots: a qualitative review of recent advances of artificial intelligence in health

Erwin Loh

**Correspondence to**

Professor Erwin Loh, Monash Centre for Health Research and Implementation, Monash University, Clayton, VIC 3168, Victoria, Australia; erwin.loh@monash.edu

Received 14 February 2018

Revised 1 May 2018

Accepted 11 May 2018

Published Online First

1 June 2018

**ABSTRACT**

Artificial intelligence (AI) has the potential to significantly transform the role of the doctor and revolutionise the practice of medicine. This qualitative review paper summarises the past 12 months of health research in AI, across different medical specialties, and discusses the current strengths as well as challenges, relating to this emerging technology. Doctors, especially those in leadership roles, need to be aware of how quickly AI is advancing in health, so that they are ready to lead the change required for its adoption by the health system. Key points: 'AI has now been shown to be as effective as humans in the diagnosis of various medical conditions, and in some cases, more effective.' When it comes to predicting suicide attempts, recent research suggest AI is better than human beings. 'AI's current strength is in its ability to learn from a large dataset and recognise patterns that can be used to diagnose conditions, putting it in direct competition with medical specialties that are involved in diagnostic tests that involve pattern recognition, such as pathology and radiology'. The current challenges in AI include legal liability and attribution of negligence when errors occur, and the ethical issues relating to patient choices. 'AI systems can also be developed with, or learn, biases, that will need to be identified and mitigated'. As doctors and health leaders, we need to start preparing the profession to be supported by, partnered with, and, in future, potentially be replaced by, AI and advanced robotics systems.

**INTRODUCTION**

Artificial intelligence (AI) has been defined by Alan Turing, the founding father of AI, as 'the science and engineering of making intelligent machines, especially intelligent computer programs'.<sup>1</sup> AI in health uses algorithms and software to approximate the cognition undertaken by human clinicians in the analysis of complex medical data. AI research has been divided into subfields, based on goals such as machine learning or deep learning, and tools such as neural networks, a subset of machine learning.<sup>2</sup> AI has the potential to significantly transform the role of the doctor and revolutionise the practice of medicine, and it is important for all doctors, in particular those in positions of leadership within the health system, to anticipate the potential changes, forecast their impact and plan strategically for the medium to long term.

The impact of automation and robotics have been felt by blue-collar jobs for a while. A recent working paper by the National Bureau of Economic Research found that the arrival of one new industrial robot

in a local labour market coincides with an employment drop of 5.6 workers.<sup>2</sup> Last year alone, there have been news reports of apple-picking robots,<sup>3</sup> burger-flipping robots<sup>4</sup> and a barista robot that makes you coffee.<sup>5</sup> *Nature* even ran an editorial on sex robots.<sup>6</sup>

There is a false sense of security in assuming that automation will only impact blue-collar type work that requires more manual, repetitive actions and less intellectual input. PwC released a report based on a survey of 2500 US consumers and business leaders, which predicts that AI will continue to make in-roads into white collar industries.<sup>7</sup> A large stockbroking firm ran a trial in Europe of its new AI program this year that showed it was much more efficient than traditional methods of buying and selling shares.<sup>8</sup> A Japanese insurance firm replaced 34 employees with an AI system, which it believes will increase productivity by 30% and see a return on its investment in less than 2 years.<sup>9</sup> The *Washington Post* used an AI reporter to publish 850 articles in the past year.<sup>10</sup>

Not even the jobs of computer programmers, the creators of the code for AI, are safe. Microsoft and Cambridge built an AI capable of writing code that would solve simple math problems.<sup>11</sup> Lawyers are not exempt either. Late last year, an AI was able to predict the judicial decisions of the European Court of Human Rights with 79% accuracy.<sup>12</sup>

Compared with other industries like hospitality or airlines, health has been a relative slow adopter of electronic systems, such as electronic health record (EHR) systems, which have only recently become mainstream.<sup>13</sup> Similarly, although AI is now embedded in many forms of technologies such as smartphones and software, its use in the front-line of clinical practice remains limited. Nevertheless, research in this area continues to grow exponentially.

**QUALITATIVE REVIEW METHODOLOGY**

This paper summarises the past 12 months of health research in AI, across different medical specialties, and discusses the current strengths and weaknesses, as well as challenges, relating to this emerging technology. The author notes that much progress has been made by AI developments in health over the past two to three decades and has focused on the past 12 months because of some of the exponential gains made, mainly due to improvements in computer hardware technologies. The author has specifically restricted his review to recent research in AI published in high-ranking peer-reviewed medical



To cite: Loh E. *BMJ Leader* 2018;**2**:59–63.

journals. The selection criteria involved keywords relating to artificial intelligence, machine learning, deep learning and algorithms relating to medical diagnosis, planning and treatment.

This qualitative review is not intended to be a systematic review, and the author has restricted the research to AI research that will likely to have the most impact to clinical practice, a judgement that is subjective to the author's own experience and expertise as a specialist medical administrator in both academia and practice. The time period of around 12 months is because the exponential growth and improvements in AI technology means that any data presented that are older may no longer be applicable.

The focus of the review is to provide a high-level update of recent AI research in health to ensure that medical practitioners, especially those in leadership roles, are made aware of how quickly AI is advancing in health, so that they are made ready to lead the change required for its adoption by the health system.

## FINDINGS

### AI in medical diagnosis

AI has now been shown to be effective in the accurate diagnosis of various medical conditions. For example, in ophthalmology, an AI-based grading algorithm was used to screen fundus photographs obtained from diabetic patients and identify, with high reliability (94% and 98% sensitivity and specificity), to determine cases that should be referred to an ophthalmologist for further evaluation and treatment.<sup>14</sup> In another study, researchers showed that an AI agent, using deep learning and neural networks, accurately diagnosed and provided treatment decisions for congenital cataracts in a multihospital clinical trial, performing just as well as individual ophthalmologists.<sup>15</sup>

In relation to skin cancer, researchers trained a neural network using a dataset of 129 450 clinical images and tested its performance against 21 board-certified dermatologists on biopsy-proven clinical images. The neural network achieved performance on par with all tested experts, demonstrating that an AI was capable of classifying skin cancer with a level of competence comparable with dermatologists.<sup>16</sup> In another study using routine clinical data of over 350 000 patients, machine learning significantly improved accuracy of cardiovascular risk prediction, correctly predicting 355 (additional 7.6%) more patients who developed cardiovascular disease compared with the established algorithm.<sup>17</sup>

Clinical neuroscience has also benefited from AI. A deep-learning algorithm used MRI of the brain of individuals 6 to 12 months old to predict the diagnosis of autism in individual high-risk children at 24 months, with a positive predictive value of 81%.<sup>18</sup> Similarly, in another study, a machine learning method designed to assess the progression to dementia within 24 months, based on a single amyloid PET scan, obtained an accuracy of 84%, outperforming the existing algorithms using the same biomarker measures and previous studies using multiple biomarker modalities.<sup>19</sup>

### AI in psychiatry

AI may be good at diagnosing physical illness, but what about its use in psychological medicine and psychiatry? The emerging literature has also shown that AI is proving to be useful in these clinical areas. For example, researchers built a predictive model based on machine learning using whole-brain functional magnetic resonance imaging (fMRI) to achieve 74% accuracy in identifying patients with more severe negative and positive symptoms in schizophrenia, suggesting the use of brain imaging

to predict the disease and its symptom severity.<sup>20</sup> In another study, researchers demonstrated that a linguistic machine learning system, using fMRI and proton magnetic resonance spectroscopy (<sup>1</sup>H-MRS) inputs, showed nearly perfect classification accuracy and was able to predict lithium response in bipolar patients with at least 88% accuracy in training and 80% accuracy in validation, allowing psychiatrists the ability to predict lithium response and avoid unnecessary treatment.<sup>21</sup>

It is one thing for AI to be able to recognise patterns on images from radiology and pathology tests. Can AI be as good as psychiatrists when it comes to predicting mental health conditions that do not have a clear biomarker? A landmark paper of a meta-analysis of 365 studies spanning 50 years published by the American Psychological Association found that prediction of suicide was only slightly better than chance for all outcomes, and that this predictive ability has not improved across 50 years of research, leading the authors to suggest the need for a shift in focus from risk factors to machine learning-based risk algorithms.<sup>22</sup>

Researchers at the Vanderbilt University Medical Centre created machine-learning algorithms that achieved 80%–90% accuracy when predicting whether someone will attempt suicide within the next 2 years, and 92% accuracy in predicting whether someone will attempt suicide within the next week, by applying machine learning to patients' EHRs. In other words, when it comes to predicting suicide attempts, AI appears to be better than human beings, although the clinical applicability in the real world remains unproven.<sup>23</sup> In another study, researchers used machine-learning algorithms to identify individuals at risk of suicide with high (91%) accuracy, based on their altered fMRI neural signatures of death-related and life-related concepts.<sup>24</sup> These developments in AI are now being applied. Facebook is one of several companies exploring ways to use AI algorithms to predict suicide based on mining social media.<sup>25</sup>

### AI in treatment

So, we have established that AI can be helpful in predicting mental health conditions, but can AI also be helpful in the provision of psychological treatments? Researchers found that soldiers are more likely to open up about post-traumatic stress when interviewed by a computer-generated automated virtual interviewer, and such virtual interviewers were found to be superior to human ones in obtaining more psychological symptoms from veterans.<sup>26</sup>

What about robot surgeons? Robotic surgical devices already exist, but they still require human control—is AI able to perform autonomous surgery without human input? In a robotic surgery breakthrough in 2016, a smart surgical robot stitched up a pig's small intestines completely on its own and was able to do a better job on the operation than human surgeons who were given the same task.<sup>27</sup> What is even more impressive is that late last year, a robot dentist in China was able to carry out the world's first successful autonomous implant surgery by fitting two new teeth into a woman's mouth without any human intervention.<sup>28</sup>

### AI's current strengths

So, based on the available evidence, what is AI good at today? It is clear that AI's current strength is in its ability to learn from a large dataset and recognise patterns that can be used to diagnose conditions. This puts AI in direct competition with medical specialties that are involved in diagnostic tests that involve pattern recognition, and the two obvious ones are pathology and radiology.

An editorial on recent studies point to the future of computational pathology, suggesting that computers will increasingly become integrated into the pathology workflow when they can improve accuracy in answering questions that are difficult for pathologists.<sup>29</sup> However, Google researchers used an AI in a study to identify malignant tumours in breast cancer images with an 89% accuracy rate, compared with 73% achieved by a human pathologist.<sup>30</sup> In another study, deep learning algorithms achieved better diagnostic performance than a panel of 11 pathologists, in a simulated time-constrained diagnostic setting, in detecting lymph node metastases in tissue sections of women with breast cancer.<sup>31</sup>

Similarly, radiologists are grappling with the potentially disruptive applications of machine learning to image analysis in their specialty, but remain as a profession optimistic that AI will be able to provide opportunities for radiologists to augment and improve the quality of care they provide to their patients.<sup>32</sup> However, AI systems continue to improve in their diagnostic and predictive capabilities in radiology. For example, a machine-learning model, using three-dimensional cardiac motion on cardiac MRI, was able to predict survival outcome independent of conventional risk factors in patients with newly diagnosed pulmonary hypertension.<sup>33</sup> It is also interesting to note that the first United States Food and Drugs Administration approval for an AI application in a clinical setting is for a deep learning platform in radiology, to help doctors diagnose heart problems.<sup>34</sup>

### Can AI completely replace the role of a doctor?

AI may be as good as, or even better than, humans when it comes to formulating diagnoses based on recognising patterns on images, but is AI ready to take over the complete role of a fully trained medical practitioner? So far, the answer appears to be—not yet. In the first direct comparison of diagnostic accuracy, physicians were found to vastly outperform computer algorithms in diagnostic accuracy (84.3% vs 51.2% correct diagnosis in the top three listed).<sup>35</sup> Bear in mind that this study compared doctors with relatively simple symptom checker applications.

In a more recent study, Watson, IBM's AI platform, took just 10 min to analyse a genome of a patient with brain cancer and suggest a treatment plan, compared with human experts who took 160 hours to make a comparable plan.<sup>36</sup> In another study, Watson found cancer treatments that oncologists overlooked, by discovering 'potential therapeutic options' for 323 additional patients after analysing 'large volumes of data', including past studies, databases and genetic information.<sup>37</sup> It should be noted that these superior performances in the theoretical setting has not translated well into real-world clinical practice, based on recent reports of poor clinician adoption at a major American cancer centre.<sup>38</sup>

As such, it would seem that AI systems may be better than human doctors in coming with diagnoses or management plans, if they are provided with sufficiently large amounts of data that are beyond what humans can manually analyse.

## DISCUSSION

### Challenges of AI in health

It is clear from the qualitative literature review that AI in health has progressed remarkably, even within the span of 12 months looked at. It is likely that much of this recent progress is due to the increasing presence of large training data sets and improvements in computer hardware, in the form of memory and computational capacity. However, there are some challenges that need to be considered as AI usage increases in healthcare. One of

the concerns that has been raised is the issue of legal liability. If a medical error occurs, who is to be held liable? A robot surgeon is not a legal entity, so should the patient sue the owner, the programmer, the manufacturer or someone else? Could an AI ever be subject to criminal liability? These AI dilemmas are not unique to health—for example, there have already been a few high-profile self-driving car accidents, some resulting in fatalities. These are some of the issues that legal experts have been grappling with that are still unresolved.<sup>39</sup>

The other issue to consider is the potential for AI to greatly reduce the number of medical errors and misdiagnoses, and therefore reduce medicolegal claims. What happens when the ability of AI surpasses that of the average doctor? If a doctor relies on the recommendation of an AI tool, which ends up being wrong, is it still the negligence of the doctor if that tool has already been proven to be more reliable than the average doctor? An argument has been put forth, although under the US legal system, to suggest that a by-product of an increased use of AI in health is that doctors will practise less defensive medicine, by ordering less unnecessary tests, because they will be relying on the recommendations of AI systems that are better diagnosticians than they are.<sup>40</sup> In fact, there may come a day that it would be considered negligent for a doctor *not* to consider the recommendation of a health AI system if that becomes the standard of care.

There is also the matter of morality and ethics with AI. The best way to illustrate this issue is by describing the classic 'trolley problem'—if you are in a trolley that is going down a track that is about to hit five workers, and you can redirect the trolley by turning it onto another track but there is one worker on it, is it morally permissible to turn the trolley to spare the lives of five workers by killing the single worker?<sup>41</sup> This dilemma is particularly pertinent to self-driving cars, as that scenario could realistically actually happen in real life—what should the self-driving car in the event of an accident do in an attempt to reduce the number of injured humans? Should the self-driving car prioritise the passengers over the pedestrians? Who gets to make these decisions? The programmer or the passenger?

Researchers have attempted to resolve this issue by suggesting that self-driving cars be equipped with what they call an 'Ethical Knob', a device enabling passengers to ethically customise their autonomous vehicles to choose between different settings corresponding to different moral approaches or principles. In this way, the AI in self-driving cars would be entrusted with implementing users' ethical choices, while manufacturers/programmers would be tasked with enabling the user's choice.<sup>42</sup> Similarly, an AI in healthcare can be provided guidance as to the moral wishes of the patient—for example, does the patient want to maximise length of life or the quality of life?

This brings us to another real issue with AI—inherent bias. AI systems can be inadvertently programmed to have bias because of the biases of the programmers or, with the development of self-learning algorithms, actually learn to be biased based on the data it is learning from. In addition, AI systems find it more difficult to generalise findings from a narrower dataset, with minor differences from a training set potentially making larger-than-intended impact on a prospective set of data, creating potential bias. A recent study demonstrated that AI can learn to have racist or sexist biases, based on word associations that are part of data it was learning from, sourced from the internet that reflected humanity's own cultural and historical biases.<sup>43</sup> Strategies to minimise and mitigate such biases will need to be in place as adoption of AI by health increases.

The last issue that needs to be considered relates to how AI uses data. In the past, EHR systems used to require that data be

properly entered into the correct categories for the right queries to be made to extract useful information. However, the advent of fuzzy logic, a form of AI, now allows for free-text unstructured text to be queried and categorised in real time to provide meaningful information.<sup>44</sup> The quality of the information extracted is still dependent on the accuracy of the data being entered, as patient-reported outcome measures may still be unreliable.<sup>45</sup> In addition, sophisticated AI systems can link disparate health data from separate databases together to form connections that may otherwise be missed.

As such, AI is now being applied to the large health data repositories because of the amount of free-text stored and also because AI, through machine learning, needs access to vast amounts of data. However, the issue of data ownership and privacy needs to be considered. A relevant case study is the recent finding by the UK's Information Commissioner that a National Health Service trust breached privacy laws by sharing patient data with Google for Google's DeepMind Streams app.<sup>46</sup> Although this app did not directly use AI, the alleged data breach demonstrates that need for the development of a data governance framework that takes into account data ownership, privacy principles, patient consent and data security.<sup>47</sup> Current privacy laws may need to be reviewed to ensure they are relevant even as social media and other large technologies like Google start using AI to commercialise the big data they have collected from their millions of users.

### Future of AI

There is no turning back from the rise of AI in all aspects of our lives. AI already resides in the smartphones that a lot of us own, in the form of smart digital assistants. But AI has progressed beyond helpful chatbots. For example, Google's AI group, Deepmind, unveiled AlphaGo, an AI that took just 3 days to master the ancient Chinese board game of Go with no human input, as reported in *Nature*.<sup>48</sup> This version of AI was able to win against its previous version (that famously beat the world champion in Go previously) 100 games to 0. More recently, AlphaZero, another AI from Google, learnt the rules of chess in 4 hours by playing against itself 44 million times and went on to beat Stockfish, a well-established chess program.<sup>49</sup>

AI researchers are already developing AI algorithms that are able to learn, grow and mature like human beings do, through self-reflection<sup>50</sup> and experiencing the world firsthand.<sup>51</sup> AI can currently analyse large amounts of data much faster than humans can using today's hardware. However, quantum computers, which may outperform the classical computers we have today by many factors, are already in development and only a few years away.<sup>52</sup> In addition, scientists have made a pioneering breakthrough by developing photonic computer chips—that use light rather than electricity—that imitate the way the brain's synapses operate, which means that computers may be able to process data at the speed of light in the near future, compared with human nerve conduction speed that is slower than electricity as it is.<sup>53</sup>

With dramatic improvements in computer software and hardware coming online, and increasing access to large datasets that are increasingly being linked together, it is no wonder that Ray Kurzweil, a Google AI expert and well-known futurist, believes that AI will surpass the brainpower of a human being by 2023 and reach what he terms 'singularity' in 2045, which is when AI will surpass the brainpower equivalent to that of all human beings combined.<sup>54</sup>

### Implications for medical leaders

Those of us who are medical leaders in healthcare, in particular, in the public health system, know that the health system is traditionally risk averse and tends to be a slower adopter of new technologies. Nevertheless, it is essential that medical leaders like us are aware of the potential impacts that new health technologies will have on the current and future health system.

As such systems are introduced into our health services, medical leaders need to ensure that there are strong and robust governance structures in place to ensure that there is appropriate review of these new technologies prior to implementation, in terms of their safety, cost-effectiveness and that staff are credentialled to use the new technologies. A data governance framework will also be required to oversee how data are managed internally, the data standards and quality expected, how data are received, how data are secured and how data are shared externally to different stakeholders, in compliance with relevant laws and regulations. An appropriate training regime should also be implemented to ensure that staff are aware of their ethical and legal responsibilities when it comes to data management, especially as it relates to the use of social media.

Medical leaders will also need to constantly scan the horizon for future developments in the field of AI, and consider future risks and opportunities, in order to plan accordingly. AI and automation will have an impact of the health workforce, and workforce planning will need to take this issue into account. The opportunities offered by AI to improve the care of patients need to be taken into account when new IT systems are introduced, in particular, where AI can assist in interrogating large amounts of health data, which may be unstructured or separated into different silos.

Medical leaders should also be aware that AI systems are not just relevant for clinical care—AI systems are increasingly being applied in the management setting. AI can be used to support, and potentially replace, the role of managers, including in health, in financial management, priority setting, resource allocation and workforce management. We will need to consider how AI can support us in our roles, now and into the future.

Lastly, medical leaders will need to be change agents and lead the change as AI transforms the healthcare system in the coming years. We will need to ensure that the patient experience and needs are always prioritised, and that compassion and kindness are not replaced by efficiencies and metrics. As leaders of clinicians, we will need to manage the anxiety of the clinical workforce through potential uncertain times, by refocussing any changes on improving patient care. Ultimately, medical leaders are still doctors, and our duty of care is to our patients.

### CONCLUSION

It is evident from this qualitative review of recent evidence that AI research in health continue to progress, and that AI is proving to be effective in most aspects of medicine, including diagnosis, planning and even treatment. As a profession, we need to have a mature discussion and debate about the legal, ethical and moral challenges of AI in health, and mitigate any potential bias that such systems may inherit from their makers.

Regardless of whether the AI singularity comes to pass or not, AI in health will continue to improve, and these improvements appear to be accelerating. There are clear challenges for the adoption of AI in health for health services, organisations and governments, and a need to develop a policy framework around this issue. As doctors and health leaders, we need to start preparing the profession to be supported by, partnered with, and, in future,

potentially be replaced by, AI and advanced robotics systems. We have an opportunity now to literally shape the development of humanity's future autonomous health providers, and we should be leaders in this space rather than passive observers.

**Contributors** EL planned, conducted and submitted the study.

**Funding** The author has not declared a specific grant for this research from any funding agency in the public, commercial or not-for-profit sectors.

**Competing interests** None declared.

**Patient consent** Not required.

**Provenance and peer review** Not commissioned; externally peer reviewed.

**Open access** This is an open access article distributed in accordance with the Creative Commons Attribution Non Commercial (CC BY-NC 4.0) license, which permits others to distribute, remix, adapt, build upon this work non-commercially, and license their derivative works on different terms, provided the original work is properly cited and the use is non-commercial. See: <http://creativecommons.org/licenses/by-nc/4.0/>

© Article author(s) or their employer(s) unless otherwise stated in the text of the article) 2018. All rights reserved. No commercial use is permitted unless otherwise expressly granted.

## REFERENCES

- Turing AM. I.—Computing machinery and intelligence. *Mind* 1950;LIX:433–60.
- Acemoglu D, Restrepo P. *Robots and jobs: evidence from US labor markets*, 2017. NBER Working Paper No. 23285.
- Scientific American. What an Apple-Picking Robot Means for the Future of Farm Workers. [Online]. 2017 <https://www.scientificamerican.com/article/what-an-apple-picking-robot-means-for-the-future-of-farm-workers/> (cited 1 Nov 2017).
- Inc. Meet Flippy the Burger-Flipping Robot (He's Hungry to Steal Your Fast-Food Job). [Online]. 2017 <https://www.inc.com/peter-economy/meet-flippy-the-burger-flipping-robot-and-hes-hung.html> (cited 1 Nov 2017).
- Wired. This Robot Barista Makes a Dang Good Latte. [Online]. 2017 <https://www.wired.com/2017/01/cafe-x-robot-barista/> (cited 1 Nov 2017).
- Nature. Let's talk about sex robots. *Nature* 2017;547:138.
- PwC. *Bot.Me: A Revolutionary Partnership*, 2017. Consumer Intelligence Series.
- Financial Times. JPMorgan develops robot to execute trades. 2017 <https://www.ft.com/content/16b8ffb6-7161-11e7-aca6-c6bd07df1a3c> (cited 1 Nov 2017).
- Guardian. Japanese company replaces office workers with artificial intelligence. [Online]. 2017 <https://www.theguardian.com/technology/2017/jan/05/japanese-company-replaces-office-workers-artificial-intelligence-ai-fukoku-mutual-life-insurance> (cited 1 Nov 2017).
- Digiday. The Washington Post's robot reporter has published 850 articles in the past year. [Online]. 2017 <https://digiday.com/media/washington-posts-robot-reporter-published-500-articles-last-year/> (cited 1 Nov 2017).
- Balog M, Gaunt A, Brockschmidt M, et al. DeepCoder: learning to write programs. *ICLR* 2017.
- Aletras N, Tsarapatsanis D, Preotjuc-Pietro D, et al. Predicting judicial decisions of the European Court of Human Rights: a Natural Language Processing perspective. *PeerJ Computer Science* 2017;2:e92.
- Palabindala V, Pamarthy A, Jonnalagadda NR. Adoption of electronic health records and barriers. *J Community Hosp Intern Med Perspect* 2016;6:32643.
- Gargeya R, Leng T. Automated identification of diabetic retinopathy using deep learning. *Ophthalmology* 2017;124:962–9.
- Long E, Lin H, Liu Z, et al. An artificial intelligence platform for the multihospital collaborative management of congenital cataracts. *Nat Biomed Eng* 2017;1:0024.
- Esteve A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 2017;542:115–8.
- Weng SF, Reps J, Kai J, et al. Can machine-learning improve cardiovascular risk prediction using routine clinical data? *PLoS One* 2017;12:e0174944.
- Hazlett HC, Gu H, Munsell BC, et al. Early brain development in infants at high risk for autism spectrum disorder. *Nature* 2017;542:348–51.
- Mathotaarachchi S, Pascoal TA, Shin M, et al. Identifying incipient dementia individuals using machine learning and amyloid imaging. *Neurobiol Aging* 2017;59:80–90.
- Gheiratmand M, Rish I, Cecchi GA, et al. Learning stable and predictive network-based patterns of schizophrenia and its clinical symptoms. *NPJ Schizophr* 2017;3:22.
- Fleck DE, Ernest N, Adler CM, et al. Prediction of lithium response in first-episode mania using the LITHium Intelligent Agent (LITHIA): pilot data and proof-of-concept. *Bipolar Disord* 2017;19:259–72.
- Franklin JC, Ribeiro JD, Fox KR, et al. Risk factors for suicidal thoughts and behaviors: a meta-analysis of 50 years of research. *Psychol Bull* 2017;143:187–232.
- Walsh CG, Ribeiro JD, Franklin JC. Predicting risk of suicide attempts over time through machine learning. *Clin Psychol Sci* 2017;5:457–69.
- Just MA, Pan L, Cherkassky VL, et al. Machine learning of neural representations of suicide and emotion concepts identifies suicidal youth. *Nat Hum Behav* 2017;1:911–9.
- Reardon S. AI algorithms to prevent suicide gain traction. *Nature* 2017;64. ISSN 1476-4687.
- Lucas GM, Rizzo A, Gratch J, et al. Reporting mental health symptoms: breaking down barriers to care with virtual human interviewers. *Frontiers in Robotics and AI* 2017;4:51.
- Shademan A, Decker RS, Opfermann JD, et al. Supervised autonomous robotic soft tissue surgery. *Sci Transl Med* 2016;8:337ra64.
- Gruth China Morning Post. Chinese robot dentist is first to fit implants in patient's mouth without any human involvement. [Online]. 2017 <http://www.scmp.com/news/china/article/2112197/chinese-robot-dentist-first-fit-implants-patients-mouth-without-any-human> (cited 10 Jan 2018).
- Granter SR, Beck AH, Papke DJ. AlphaGo, Deep Learning, and the Future of the Human Microscopist. *Arch Pathol Lab Med* 2017;141:619–21.
- Liu Y, Gadepalli K, Norouzi M, et al. Detecting cancer metastases on gigapixel pathology images. *arXiv* 2017:1703.
- Ehteshami Bejnordi B, Veta M, Johannes van Diest P, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA* 2017;318:2199–210.
- Kruskal JB, Berkowitz S, Geis JR, et al. Big Data and Machine Learning—Strategies for Driving This Bus: A Summary of the 2016 Intersociety Summer Conference. *J Am Coll Radiol* 2017;14:811–7.
- Dawes TJW, de Marvao A, Shi W, et al. Machine learning of three-dimensional right ventricular motion enables outcome prediction in pulmonary hypertension: a cardiac MR imaging study. *Radiology* 2017;283:381–90.
- Forbes. First FDA Approval For Clinical Cloud-Based Deep Learning In Healthcare. [Online]. 2017 <https://www.forbes.com/sites/bernardmarr/2017/01/20/first-fda-approval-for-clinical-cloud-based-deep-learning-in-healthcare/> (cited 11 January 2018).
- Semigran HL, Levine DM, Nundy S, et al. Comparison of physician and computer diagnostic accuracy. *JAMA Intern Med* 2016;176:1860–1.
- Wrzeszczynski KO, Frank MO, Koyama T, et al. Comparing sequencing assays and human-machine analyses in actionable genomics for glioblastoma. *Neurol Genet* 2017;3:e164.
- Patel NM, Michelini VV, Snell JM, et al. Enhancing next-generation sequencing-guided cancer care through cognitive computing. *Oncologist* 2018;23:179–85.
- Stat News. IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close. [Online]. 2017 <https://www.statnews.com/2017/09/05/watson-ibm-cancer/> (cited 3 Apr 2018).
- Kingston J. *Artificial Intelligence and Legal Liability. Research and Development in Intelligent Systems XXXIII*. Cham, Switzerland: Springer, 2016.
- Thomas S. *Artificial Intelligence, Medical Malpractice, and the End of Defensive Medicine*, 2017.
- Thomson JJ. The trolley problem. *Yale Law J* 1985;94:1395–415.
- Contissa G, Lagioia F, Sartor G. The Ethical Knob: ethically-customisable automated vehicles and the law. *Artificial Intelligence and Law* 2017;25:365–78.
- Caliskan A, Bryson JJ, Narayanan A. Semantics derived automatically from language corpora contain human-like biases. *Science* 2017;356:183–6.
- Yager RR. Fuzzy logics and artificial intelligence. *Fuzzy Sets Syst* 1997;90:193–8.
- Frost MH, Reeve BB, Liepa AM, et al. What is sufficient evidence for the reliability and validity of patient-reported outcome measures? *Value Health* 2007;10:S94–105.
- BBC News. Google DeepMind NHS app test broke UK privacy law. [Online]. 2017 <http://www.bbc.com/news/technology-40483202> (cited 3 Apr 2018).
- Commissioners 3ICoDPa. *Artificial Intelligence, Robotics, Privacy and Data Protection*. Marrakech: European Data Protection Supervisor, 2016:3. In Room document.
- Silver D, Schrittwieser J, Simonyan K, et al. Mastering the game of Go without human knowledge. *Nature* 2017;550:354–9.
- Silver D, Hubert T, Schrittwieser J, et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv* 2017.
- Jaderberg M, Mnih V, Czarnecki WM, et al. Reinforcement learning with unsupervised auxiliary tasks. *arXiv* 2016.
- Denil M, Agrawal P, Kulkarni T, et al. Learning to perform physics experiments via deep reinforcement learning. *arXiv* 2017:1611.01843.
- Zhang J, Pagano G, Hess PW, et al. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature* 2017;551:601–4.
- Cheng Z, Rios C, Pernice WHP, et al. On-chip photonic synapse. *Sci Adv* 2017;3:e1700160.
- Kurzweil R. *The Singularity Is Near: When Humans Transcend Biology*. New York: Penguin, 2006.