Artificial intelligence and physicians in the future of medicine: a meeting of minds?

Graham McLeod BSc*, Jagandeep Brar BSc

*Max Rady College of Medicine, University of Manitoba
727 McDermot Avenue, Winnipeg, R3E 3P5

Abstract

Artificial intelligence (AI) is one of today's most powerful technologies. Having already transformed the business world, AI may be poised to transform healthcare next. Current AI systems demonstrate impressive competency in certain tasks of clinical medicine. Machine learning approaches to creating AI are of particular relevance in healthcare, given the ability of modern machine learning algorithms to work with large amounts of complex data and generate intelligent predictions therefrom. Here we propose that much of what physicians do can be modelled as information processing and thus can be performed by AI. We further propose that whereas certain AI systems may adopt approaches based on novel pattern extraction and interpretation, and thus diverge from human physician cognition, AI is well-positioned to assist physicians by operating in parallel alongside them. Navigating the intersection of physician and AI competence will be a tremendous and complex challenge, but may return high rewards in improving patient outcomes and lead to transformative gains in medical knowledge. Advances in AI will have tremendous and complex impact on the future of medicine.

Keywords: artificial intelligence, machine learning, information theory, computer science, healthcare, medicine

1. Introduction

Artificial intelligence (AI) may be the most transformative technology of the 21st century. Any non-human machine system performing intelligent behaviour — behaviour that is proficient with respect to a complex goal — falls under the rubric of AI.[1] Recent years have brought tremendous advances in AI, with certain AI systems now capable of human-level speech recognition,[2,3] human-level language translation,[4] superhuman image recognition,[5] and superhuman performance in numerous complex games such as poker,[6,7] Go,[8] and Capture the Flag.[9]

Alongside development of enhanced capabilities, use of and interest in AI is also growing. The 2018 AI Index Report, prepared by the Human-Centered AI Institute at Stanford University, documents manifold aspects of progress, including huge increase in the number of AI papers published per year, increased number of AI startups and patents, and growing widespread adoption of AI in industry.[10] In some sectors, as of 2018, as many as 75% of companies had trialed AI or were currently using AI for certain functions of their business.[10] AI is currently used by diverse collections of companies in various industries including technology (Google, Samsung, Apple), social media (Twitter, Facebook, Instagram), entertainment media (Spotify, Netflix, Walt Disney), consumer goods (Amazon, Walmart), food and beverage (e.g. Starbucks, Coca-Cola, McDonald's), transportation (Hopper, Uber), and automotive (Tesla, BMW, Volvo).[11]

In view of increasing capabilities of AI systems and simultaneous transformative benefits of AI in diverse industries, many have suggested that AI may next transform healthcare.[12,13,14,15,16] Specialized AI systems have already been deployed into healthcare in various regions worldwide. In the US, the FDA has already approved specific AI algorithms for tasks including interpretation of magnetic resonance (MR) and computed tomography (CT) images of the brain,[17] heart,[18] liver, and lungs,[19] with some degree of autonomy.[20,21] Moreover, the FDA is fast-tracking approval of further AI algorithms.[22] In Japan, IBM's Watson now assists in diagnosing leukemia via genome sequencing, with some notable success, including possibly having saved lives.[23] In India, AI is widely used to interpret urgent ECGs to rule out or diagnose myocardial infarctions[24] and to detect cervical cancer in pathologic samples.[25] One hospital in Guangzhou, China is using AI to suggest diagnoses for hundreds of diseases, interpret computed tomography (CT) scans, and organize patient files via facial recognition.[26] Meanwhile, software by Beijing-based AI start-up, Infervision, assists with interpretation of CT scans in a majority of Chinese hospitals, as

*correspondence to: mcleodg5@myumanitoba.ca

well as select additional hospitals outside of China.[27] In Europe, Google DeepMind is partnered with UK's National Health Service as of 2015, and has since been receiving patient data for the purpose of developing AI systems for healthcare.[28] At St. Michael's Hospital in Toronto, Canada, the Vector Research Institute is testing the application of AI to improve interpretations of radiographic imaging studies[29] and to predict when hospitalized patients require transfer to the intensive care unit.[30] These initial examples form what is likely just the beginning of AI's involvement in healthcare.

Within the roles of clinical medicine, the competencies of AI seem particularly well-suited to certain tasks, such as image classification. Much has been discussed about the possibility of AI replacing diagnostic radiologists.[31,32,33,34] Notably, many of the above examples of AI's deployment into healthcare are in diagnostic imaging. However, deployments for broader purposes have already begun and are likely to continue. As we will see, given extreme broadness of certain principles of AI and of computation more generally, the growing competency of AI applications may expand to cover extensive aspects of clinical medicine. Recent research indicates AI competence in diverse tasks such as predicting treatment response to medications,[35,36] predicting cardiovascular risk from routine laboratory data,[37] and diagnosing rheumatoid arthritis via automatic analysis of patients' electronic health records.[38]

Given broad and increasing capabilities of medical AI, some have wondered if AI of the future could render human physicians obsolete. Silicon Valley investor Vinod Khosla suggests that AI will replace 80% of doctors, possibly even on the timescale of a couple decades (written in 2012).[39] Notably, an AI-powered robot named Xiaoyi recently passed the Chinese medical licensing exam.[40] Although it is possible that advanced future AI systems will someday usurp certain roles of physicians, healthcare will first be faced with increasing competency of AI in an increasing number of medical applications, with specialized medical AI systems performing at levels of competency approaching, equalizing, and/or ultimately surpassing the competency of human physicians. Thus, rather than focusing attention on the possibility of physician replacement by AI, the more pressing question of today is *"what should the healthcare system do with AI systems that are as good as or better than physicians?"*

To adequately discuss these crucial matters, and to be prepared for what future technological advances will bring, physicians and healthcare policymakers must attain a working knowledge of the possibilities of AI in medicine. This will require some understanding of computational principles. Despite increasing adoption rates and ubiquity, AI and computers can be highly counterintuitive. For example, researchers at the University of Wisconsin-Madison recently created an AI system in a "nanophotonic medium" — essentially, a piece of glass, with no electricity required for its operations — that classified handwritten digits.[41] Even within standard modern electronic computers, AI is a diverse and immense category, subsections of which may defy intuitions. Thus, understanding AI's conceptual underpinnings in terms of computer science and information theory is prerequisite to understanding possible broad futures of AI in medicine.

In this paper we discuss physical and philosophical underpinnings of how AI systems may achieve competence in the complex goals of healthcare, noting certain similarities between information processing performed by AI and by human physicians. We also discuss certain capabilities of medically-focused AI systems, highlighting multi-faceted ways in which AI and physician competence will intersect, and how this will benefit patient outcomes and advance medical knowledge. To make the most of the future, healthcare must place a high priority on capturing value and mitigating risks of AI in the future of medicine.

2. Computation and artificial intelligence: brief primer for a medical audience

This section describes fundamentals of computation and AI in order to see how they may be relevant to medicine. We will see that computers are physical systems that store and transform information, and, notably, how this abstract framework can map much of what occurs in medicine. Transforming information in certain ways earns a computational system the designation of "intelligence", and machine learning approaches are one effective route by which contemporary computing systems achieve intelligence. We will also discuss certain advantages conferred by different types of AI. Readers familiar with computation and AI should advance to Section 3.

2.1 Information: modeling the world

A computer is a mechanical device that stores and transforms information via physical processes. Thus, the extent to which realms of medicine can be accessed by computers and AI is wholly dependent on the extent to which medicine can be modelled in terms of information. We will consider information generally before advancing to specifically consider information in medicine in Section 3.

Information has various technical definitions relating to divergent subcultures in the philosophy of information.[42] For our purposes, a sufficiently philosophically-neutral definition will be "what is conveyed or represented by a particular arrangement or sequence of things".[43] For example, a geographic map carries information about the physical environment of earth, in that multicolored patterns of shapes and their word-labels (the particular arrangement of things) convey and represent the earth. In carrying information about the earth (or a portion thereof), it can be said that the map *models* the earth (or its portion thereof). Whereas a map models the world via apparent physical likeness (e.g. water on a map may be coloured blue, mountains may appear raised), other technologies achieve modelling by means that are highly abstract.

For example, certain modern technologies encode information via etching microscopic pits into CDs, magnetizing surface points on computer hard drives, and using electrons to influence the charge of a capacitor.[44] Although it may be counterintuitive, these highly abstract modes of instantiating information can render extremely faithful models of systems in the outside world. A map stored on a CD or in a computer hard drive can model the outside system of interest (in this case, a particular geographic region) with the same arbitrary closeness that can be achieved by a paper map (i.e. bounded only by storage capacity, namely the size of the paper or the size of the hard drive's memory).

As another example of abstract information storage, consider a photographic image. It can be stored digitally on a smartphone device and represented in solid-state storage (a type of computer memory that uses electrical circuits and lacks moving parts), stored on a desktop computer represented in hard disk drives (a type of computer memory utilizing spinning electromagnetic rotating disks), or printed physically onto a piece of paper and represented via molecules of ink. Some "likeness" of the image, expressed via abstract electrical, magnetic, and molecular patterns, is translated across mediums, despite distinct modes of physical instantiation. The information "has a life of its own."[44] The ability of information to flow from model to model in this way may appear peculiar, but information itself ("that which is conveyed or represented") is *substrate-independent*, meaning that it does not change according to the way in which it is stored.[45] Another interesting property of information is that there is no apparent physical constraint on what can be represented by information.[46] Thus, to model any system, one need only possess information about that system, and possess a means of instantiating that information into a device such as a computer. As a result of these properties, information stored in a computer can represent diverse and complex features of the world such as images,[47] earthquakes,[48] and quasars.[49]

2.2 Computation: transforming the model

Although all computers store information, this alone is not sufficient to achieve their designation as computer. A *computer* is a device that not only stores information in memory, but utilizes functions to transform information.[44] *Functions* are mathematical equations that accept a set of data and output a paired set of data in one-to-one correspondence. Essentially, information enters a function, is acted upon by the function, and emerges transformed. The particular transformation that occurs is specified by the particular function. For example, a function may transform its input information by a two-fold factor of multiplication, as in the simple function "y = 2x". But functions can also perform transformations of much greater complexity. Physicist Max Tegmark gives the further examples of a function transforming input information that represents current positions of chess pieces on a chessboard into information representing best next move for Black, or transforming information representing all the world's financial data into lucrative stock market purchases.[44] (Functions of this kind illustrate, in outline, how a computational algorithm can be "intelligent." Section 2.3 considers this in detail.) The process of implementing a series of arithmetic functions alongside non-arithmetic functions to retrieve an output set of data is called an *algorithm*. The process of implementing such algorithms are called *computations*; thus, circularly, a *computer* designates a device that implements algorithms.

Like information, computation is also substrate-independent. What matters for computation is the transformation of information, not the physical substrate that implements the transformation. Thus, a wide variety of mechanical systems can function as computers. This insight, in general form, appears to have been first arrived at by Spanish polymath Ramon Llull (deceased 1316), who realized that mechanical artifacts could perform "useful reasoning".[1] Thereafter, in the 1500-1700s, various simple devices were built to perform mathematical calculations.[1] Modern computers use *bits*, simple two-state storage devices, to store and transform information. Like a power switch, a single bit can be off (represented as 0), or on (represented as 1), but not in-between. ("On" and "off" correspond to mutually exclusive physical states, such as whether electrical current flows through a given wire or not, whether a given area is magnetized or not, etc.) With a large enough number of bits stored in large and complex arrays, *any amount* of information can be stored and transformed in a computer's memory,[50] giving modern computers tremendous reach.

Modern computing is generally agreed to have begun with Alan Turing's seminal 1936 paper, "On Computable Numbers, with an Application to the Entscheidungs problem", in which Turing demonstrates that if a computer can perform a minimum set of basic operations, then it is a *universal computer*, meaning it can compute anything that any other computer can compute.[1,51] Notabbly, modern electronic computers, including smartphones and laptops, are universal.[44] Given that a computer is simply a physical system that transforms information, and given the possibility of information to represent any complex and interesting feature of the world, there is no obvious limit to the kinds of interesting and useful transformations of information that potential future computers can accomplish, other than the limits imposed by the laws of physics themselves — i.e. laws regarding the kinds of systems that can instantiate information and can perform a given physical transformation, how fast a given transformation can be performed, etc.[52] In creating universal computers, humanity may have initiated what physicist David Deutsch calls "a beginning of infinity".[46]

To whatever extent the conceptual future of computation is limitless, likewise, the application of computation to solve problems in medicine may be correspondingly limitless. The next section will discuss specifically how certain computational systems achieve "intelligence".

2.3 Intelligent computation in machines and in medicine

How can a computer achieve intelligence? The answer may be implied by the previous sections, but is worth elaborating in further detail. *Intelligence* has no single standard definition, but in regards to AI may be thought of as proficiency with respect to a complex goal.[44] (What exactly constitutes "complex" is itself difficult to define. A working definition put forth by biologist Richard Dawkins proposes that something is complex if it has "some quality, specifiable in advance, that is highly unlikely to have been acquired by random chance alone".[53]) A complex goal might then be "one that is unlikely to be reached by chance alone." An example of a complex goal might be winning a game of chess. If an AI system receives input information representing the positions of chess pieces on a chess board and successfully outputs moves for Black that are better than a random move generator, then it may be said to have some degree of intelligence with respect to chess. The better the moves, the more intelligent the system. Likewise, if an AI system receives input information representing all the world's financial data and successfully outputs lucrative stock market purchases, then it may be said to have some degree of intelligence with respect to the stock market.

In addition to the complex goals of winning chess or succeeding on the stock market, AI systems can also be intelligent with respect to the goals of medicine. If an AI system receives input of pixels representing a chest x-ray and outputs correct diagnoses therefrom, the system is intelligent with respect to interpreting chest x-rays. If an AI system receives input representing vital signs and bloodwork results and outputs accurate predictions of two-week mortality therefrom, the system is intelligent with respect to predicting mortality. If an AI system receives input representing history of presenting illness, physical exam findings, and laboratory data, and outputs a correct diagnosis therefrom, the system is intelligent with respect to diagnostics.

In their capacity for modeling and transforming information about the world, computers may succeed with respect to a complex goal, medical or otherwise, and thus achieve intelligence.

2.4 Machine learning to achieve artificial intelligence

We have seen that a computational system will be designated as AI if it has the proper algorithms to transform information in a way that is of benefit towards a complex goal. But where do these intelligent algorithms come from? Early approaches to AI entailed manually inputting pre-determined rules that would enable intelligent computation.[1] However, recent successful approaches to achieving AI mostly capitalize on machine learning principles. In machine learning, a computational algorithm is designed in a way that allows it to acquire intelligent behaviour – essentially, to learn. Compared to manually inputting putatively intelligent algorithms, machine learning has generally been a more efficient approach to achieving AI. A popular and effective form of machine learning is *deep learning*. Many applications of AI in medicine are AIs that are the product of deep learning. Deep learning algorithms excel at working with complex information, uncovering useful patterns hidden within the data. Yoshua Bengio of the Université de Montréal recently gave the following description:[54]

> "Deep learning algorithms seek to exploit the unknown structure in the input distribution in order to discover good representations, often at multiple levels, with higher-level learned features defined in terms of lower-level features."

Deep learning systems are intelligent with respect to the complex goal of generating good representations of complex data. A "good representation" is one that closely maps the topography of the input data, even when the topography is very complex. Good representation tames the complexity of the input data, allowing useful patterns to be discovered in the data, and allowing the system's output to be relatively simple and easy to work with, yet faithfully representing the original complex input data. In taming complexity to arrive at simple outputs, the system may acquire intelligence with respect to certain goals pertaining to the complex input information it received.

Deep learning generally occurs in *neural networks*, a computational strategy modelled after biological brains, in which information is passed through several successive layers of computational "neurons", with each neuron transforming its information via a function. The first layer of neurons receives the input data, generally representing the data exactly as received, with the number of neurons in the first layer usually in a one-to-one ratio with the number of variables in the input data (e.g. in image classification, this may be a one-to-one neuron-to-pixel ratio). Beyond that, information is passed to a smaller number of neurons at each successive layer, each layer performing computations, mapping "features" of the data as it passes through. A given layer processes in parallel; successive layers process in series. These feature maps pass through layers until ultimately, in the final layer of neurons, the system outputs a final pattern of information. That is the general scheme, but some neural networks employ more complicated connectomes. For example, a convolutional network generates successively smaller layers until it arrives at a small enough subset to output a result. For excellent comprehensive review of neural networks, please see LeCun et al. 2015.[55]

In general, deep learning systems refine their representations on "training sets" of data, datasets in which all outputs are already known and given to the system. (For example, a deep learning system for classifying images as "cat" versus "dog" will be trained on a large number of images of cats and dogs, both labelled as such. These labels are the "ground truths" of the training stage.) The system then moves to a "validation set", in which the deep learning system creators know the output, but the system does not. This step verifies the system's ability to output a correct result within some satisfactory margin of error. Finally, the system graduates to encounter the "test set" — the data of true interest, where the output results are unknown to both the system and its creators.

How do deep learning systems learn to adopt good representations? During training, in order to successfully transform input information into the desired (known) output information, individual neurons are empowered to adjust the functions they implement via a sophisticated mathematical technique known as backpropagation. Backpropagation achieves automatic error correction by moving backwards through the deep learning algorithm to adjust the system parameters to best suit the use-case.[56] This technique essentially reverse engineers how to represent the input data across layers so as to ultimately reach the desired (known) output. Returning to the example of classifying images as cat versus dog, such a deep learning system would receive input of pixels representing cats or dogs, and from these pixelated inputs, would be tasked with generating one of two output states: one representing "cat", the other representing "dog". In training, pixelated cat images (input) are paired with the known labelled output "cat". By backpropagation, the deep learning system is slowly taught to refine its representation of cat pixels such that, with increasing success, cat pixels will flow through layers of computation in such a way that ultimately and naturally leads to the output "cat" (and likewise for dogs and dog pixels). In this way, deep learning systems are designed to learn from the data themselves, thus they are well-suited to complex information.

The above training framework, in which training set materials are labelled with ground truths, is known as *supervised learning*. The counterpart to supervised learning is *unsupervised learning*, in which the deep learning system is not provided with ground truth outputs to serve as desired endpoints. Perhaps surprisingly, even without supervised human guidance, well-designed deep learning systems can achieve very good representations of complex data, reaching high levels of competence in performing challenging feats. For example, in 2017, AlphaGo Zero learned to play Go after starting as a "blank slate" without any human data or human knowledge.8 The system did not receive ground truth labels constituting winning strategies, high score optimization, or other direct supervision; it was only provided with the rules of the game. Learning only from unsupervised iterated self-play, AlphaGo Zero went on to become the top Go player in the world, better than any human player, and defeating the previous reigning-champion AI system by a score of 100-0.[8] As deep learning systems with greater and greater learning proficiency are engineered, the complexity of problems they can solve — with minimal subsequent human input — may increase to tremendous heights.

2.5 Advantages of intelligent computers

We have defined an AI system as any computational system that transforms information in such a way as to be of use towards a complex goal. Various AI systems collectively allow at least three great general advantages to their users. First, AI may tend to achieve superhuman proficiency at certain types of information processing, such as implementing mathematical transformations quickly and reliably (as in a pocket calculator). For such tasks, a computer can function as an external cognitive prosthetic device. The scope of this advantage is widened in modern electronic computers, which can perform wide-ranging tasks quickly and reliably.

Second, by modelling the outside world (or a portion thereof), computers can provide valuable predictions about likely future events in the modelled system, such as NASA's EO-1 satellite presciently alerting human researchers of natural events before they themselves detected anything.[57] Predictive power of computational models has led to their widespread adoption in other fields such as geology (predicting earthquakes)[48] and space exploration (trajectory and payload optimizations)[58].

Third, computational models themselves can be interesting objects of study, especially those arising from unsupervised deep learning methods, which may generate models that are conceptually divergent from human knowledge and human-preprogrammed models. Thus, the AI resulting from unsupervised deep learning systems can appear foreign and other-worldly to humans. Moves made by AlphaGo Zero, the Go-winning AI generated by unsupervised machine learning algorithms, were described by champion Go players as "alien" and "from an alternate dimension".[59] Alien or not, to the extent an AI system achieves real-world success, its model likely contains representations of real-world variables and parameters that are of interest. Therefore, in examining the model of Go within which AlphaGo Zero derived its alien hyper-successful moves, a path of new insight towards the game of Go may be charted. Further, the above-mentioned EO-1 satellite AI, which alerted human scientists of events of which they had been up to then unaware, was also generated by unsupervised machine learning approaches.[57] In studying AI models such as these, we may gain new understandings of the modelled systems themselves.

The extent to which the advantages of AI will be useful in medicine depends on the extent to which that which is useful about medicine can be abstracted in terms of information and its transformations. Section 3 will briefly consider the intersection between the transformation of information and the complex goals of healthcare.

3. Information processing is central to clinical medicine

AI transforms information in ways deemed intelligent with respect to a complex goal. To what extent does this pertain to medicine?

All medicine deals with information about a patient's body. From diagnosing a disease, to recommending a medication, to forecasting the likelihood of a particular outcome within the next ten years, most of what physicians do relates to an abstract information state of their patients' bodies. Ultimately, all useful medicine is useful only in so far as it relates to information about patients' bodies. Put the other way round, there may be little use in a physician whose medical advisements *do not* correspond to any information about his or her patients' bodies.

Certain aspects of medicine such as laboratory investigations, diagnostic imaging studies, and electrophysiological studies deal directly and overtly with bodily information. Referring specifically to radiologists and pathologists, Jha and Topol went as far as to coin the term "information specialist", writing:[60]

> The primary purpose of radiologists is the provision of medical information; the image is only a means to information. Radiologists are more aptly considered "information specialists" specializing in medical imaging. This is similar to pathologists, who are also information specialists. Pathologists and radiologists are fundamentally similar because both extract medical information from images.

We propose that, in a broader and more abstract sense, all physicians are information specialists. Across medical disciplines, it may be argued that the value of physicians is to receive complex input information, process and sort it into patterns that are meaningful and actionable, and prescribe an appropriate action on the patient's behalf. Jha and Topol's information specialists provide the clearest examples — a radiologist transforms two-dimensional pixelated greyscale information into a radiographic diagnosis; a pathologist transforms complex microscopic histological information into a pathologic diagnosis. But much else of what physicians do also relates to transforming information. A typical clinical physician receives complex input — history of entrance complaint, physical examination findings, current medications, laboratory findings, other investigative findings — recognizes patterns amidst the complexity, and thereby produces comparatively simple outputs such as diagnoses, risk assessments, and prescriptions. This scenario can be modelled as information flowing through the physician and emerging transformed. There is no obvious reason why a model accomplishing an identical transformation of information cannot be programmed into a modern electronic computer. (Given the substrate-independence of information and computation, this will be possible even with modern electronic computers using computational strategies that are dissimilar to the activities of the brains of human physician.) But beyond merely modelling such a transformation of information performed in vivo by a human physician, is it conceivable that, provided with sufficient memory and processing power, a well-designed deep learning neural network could be trained to accomplish a large set of similar such transformations, much like a physician, for a large set of incoming potential patients? This would be an extraordinarily challenging feat; nonetheless, we cannot identify a law of physics or principle of computer science that would preclude this possibility. As we will see in Section 4, many pieces of the above scenario — interpreting laboratory findings, interpreting diagnostic imaging, interpreting broad clinical data as stored in electronic medical records (EMRs), etc. — have already been captured by AI models. Thus, it appears that physicians and AI alike can bring their competencies to bear on the same problems within medicine. In this general sense, physicians and AI may undergo a "meeting of minds".

Interestingly, research indicates that physician diagnostic facility depends on pattern recognition to a far greater extent than it depends on systematic reasoning from first principles.[61] Thus, both physicians and certain AI systems such as deep learning neural networks share the common feature of undergoing training on large sets of data to hone their pattern recognition abilities. In the case of physicians, data comes in the form of patients, both real (in hospital) and hypothetical (on examinations, in study materials, and human actors serving as standardized patients). Whatever patterns are recognized by physicians can likely be recognized by AI, too. Additionally, as discussed above in Sections 2.4 and 2.5, certain AI systems (in particular, those utilizing unsupervised machine learning) excel at uncovering hidden patterns. We can likely expect such systems to uncover new patterns hidden in medical information, which may have exciting implications for the future of medical knowledge. This possibility is explored further in Section 4.3.

So far we have focused heavily on the scientific aspects of medicine, discussing the extent to which these can be modelled. Narrative, humanistic, artistic aspects of medicine must be considered separately. We leave this topic as a future direction for further discussion elsewhere.

4. AI and physicians: a meeting of minds?

We have seen that AI and physicians share certain commonalities. AI transforms information in ways that are intelligent; doctors transform medical information in ways that are intelligent. Medical information can be abstractly modelled and therefore instantiated into computers, where it can easily be made available to AI systems for transformation. Therefore, from first principles, we should expect that AI can intelligently transform medical information. Further, both physicians and deep learning systems achieve their utility via a heavy reliance on pattern recognition. Deep learning, and AI generally, are active areas of research. Sophisticated machine learning algorithms of the near future, able to self-learn from staggeringly complex data and programmed into universal computers with large storage capacity and processing power, may have no near limit to the competence they can achieve. Taken collectively, these considerations persuade us that the potential for AI to succeed in the future of medicine is likely broad and extensive.

The differences between AI and physicians will be equally important as similarities. Certain AI systems are likely to form representations based on novel pattern extraction and interpretation, thus diverging from human physician cognition and models. Additionally, differences in the characteristics of models used by physicians and AI may instate accordingly different competencies across medical use cases, which may in fact be complementary. For example, AI's pixel-by-pixel analysis of diagnostic imaging studies will have higher acuity for adjacent shades of grey,[16] complementary to human physicians' generally superior capabilities of lateral thinking and broad differential diagnoses.[62,63] Thus, physicians and AI may be well-suited to augment one another in collaborative clinical practice. Collaborating to interpret diagnostic investigations and offer clinical predictions is a kind of "meeting of minds" that is already underway. Sections 4.1 and 4.2 further explore possibilities in this area.

Further, in as much as AI's models of medicine are conceptually different but nonetheless accurate and useful, examining the details of the AI models may chart a course to new insight about the body. AI systems may float free of historical biases and schemas in medicine — for example, if they begin as an agnostic deep-learning neural network, as did AlphaGo Zero. If this is the case, examining the representational models employed by medical AI systems may reveal new insights about intra-body phenomena, leading the way to paradigm shifts in medical knowledge and allowing discontinuous, transformative, and rapid advancement. The possibility that physicians may learn about medicine from AI is a second way in which a "meeting of minds" may occur. This possibility is further explored in section 4.3

The intersection of AI and physician competence is clearly complex and multi-faceted. Thus, the question "what to do with medical AI that is as good or better than doctors?" cannot be answered without specifically considering what medical AI can in fact accomplish. Having advanced a detailed physical and philosophical argument in Sections 2 and 3 that the future reach of AI in medicine may be quite broad, we will now turn attention to discussing the *current possible reach* of AI in medicine, as demonstrated by recent active research.

4.1 AI can collaborate with physicians in the information specialties

Certain tasks in medicine have received plentiful attention of early AI research. Given specific competence of certain AI systems in recognizing and classifying images,[47] a natural early step for AI in medicine has been recognizing and classifying diagnostic images. The benefits of accurate image classification will be significant. For example, if AI can successfully classify diagnostic imaging studies, it may be able to serve as a cognitive prosthetic for radiologists, leading to gains in accuracy and efficiency. Alternatively, in geographic regions underserved by radiologists, some AI systems may be accurate enough to act in lieu of radiologists. The authors of a study reporting a deep learning system for classifying pulmonary tuberculosis on chest x-ray noted that such a system could be of particular value given a relative paucity of radiologists in certain TB-endemic areas of the world.[64] Many AI systems have demonstrated competency in transforming medical imaging information, pathological information, and electrophysiological information into accurate diagnostic and predictive information. We will consider representative examples in the information specialties of AI systems with capabilities to augment or enhance the abilities of physicians.

Chest x-ray is a widely-used imaging modality that has received a plenitude of AI research attention early on. Various AI systems have demonstrated the ability to perform similarly to radiologists at interpreting chest x-rays. In a recent notable study, a deep learning system was trained to detect fourteen different pathologies as demonstrated on chest x-ray.[65] Following training on over 100,000 disease-labelled chest-x-rays, the system was found to perform radiologist-level or better on classifying eleven of fourteen pathologies. Other work has also specifically demonstrated efficacy of an AI radiologist collaboration. In a multi-centre collaboration based out of South Korea which developed a deep learning system for classifying chest x-rays with varying degrees of accuracy for various pathologies; when assisted by the deep learning system, radiologists benefited from a significant increase in sensitivity.63 The authors suggested this may be due to the AI alerting radiologists to the possibility of the presence of major thoracic disease, and to localizing the area of possible lesions to mark spots needing further attention by radiologists.[63] Researchers at Thomas Jefferson University created a radiologist-AI collaboration achieving 97.3% sensitivity and 100% specificity in classifying pulmonary tuberculosis on chest x-ray.[64] AI-radiologist collaborations may be a natural fit due to AI's pixel-by-pixel analysis allowing computation thorough approach greater acuity for adjacent shades of grey,[16] but relative inferiority at lateral thinking62 and generating a differential diagnosis.[63]

Beyond chest x-rays, deep learning systems have demonstrated early success on other imaging studies as well. A collaboration based in Australia recently reported a deep learning system that, following training and validation on nearly 50,000 frontal pelvis x-rays, achieved 97% accuracy in diagnosing hip fractures.[66] Research conducted in Budapest demonstrates a deep learning system for diagnosis breast cancer on mammography that achieved 90% sensitivity and 70% specificity (a performance considered to be on par with some physicians, but not necessarily as accurate as specialized radiologists).[67] Researchers at the University of California developed a deep learning system for classifying echocardiogram views that achieved 97.8% accuracy of classification, and when tested against electrocardiographers on single low-resolution images, achieved 91.7% accuracy versus the electrocardiographers' 70.2-84.0% accuracy.[68] Notably, the US's FDA has already approved AI algorithms for interpreting certain diagnostic imaging studies,[20] including analyzing heart hemodynamics via cardiac MR images,[18] interpreting hyperacute stroke CT brain images,[17] and evaluating liver and lung lesions evident on MR and CT images.[19]

Outside diagnostic imaging, image classifying AI systems have also achieved success in histopathological examinations. Cancer is a common indication for histopathological investigation, thus many studies have concentrated here. Google recently reported the development of "LYmph Node Assistant" (LYNA), which outperformed pathologists in diagnosing metastatic breast cancer from pathological samples,[69] and, when applied to assist pathologists, made their job of diagnosing breast cancer "easier".[70] Google also recently reported a deep-learning system which outperformed general pathologists at grading prostate cancer, achieving an overall accuracy of 70% compared the pathologists' average accuracy of 61%.[71]

Image classification techniques and other deep learning systems may also be of use in interpreting "images" of electrophysiologic tracings such as ECGs and EEGs. Regarding ECG, evidence demonstrates utility of AI to interpret ECGs to identify arrhythmias and cardiac contractile dysfunction with approximately the same accuracy as cardiologists.[72,73,74] Regarding EEG, in long-term ambulatory EEG-monitored patients, AI may be able to predict seizure onset. A recent study reports a deep learning system achieving seizure prediction accuracy of 99.6% with a prediction time of one hour ore-ictal, and a low false alarm rate of one false alarm generated every 250 hours.[75] Applied to invasive intracranial EEG, deep learning systems may have some utility in helping classify seizure onset zone.[76] Additionally, there is promising evidence regarding the utility of AI systems to monitor against seizures in the intensive care unit, with accuracy approximately as good as electroencephalographers, in less time needed to review the EEGs.[77]

These above works demonstrate AI performing certain tasks at levels of competence similar to radiologists, pathologists, and electrophysiologists. On these narrow, information-heavy tasks, the "meeting of minds" has already begun. AI in cooperation with radiologists may improve sensitivity and accuracy,[63,64] and may provide similar such advantages to pathologists.[70] In collaboration with information specialists, AI could take many roles, perhaps screening all images and alerting physicians to the likely presence of major disease, or perhaps providing a useful consultant second opinion on an as-needed basis. As noted, AI may see things differently than humans, conferring an advantage in the spirit of "two sets of eyes are better than one".

Faster interpretation times achieved by AI will be of specific use as well. The FDA-approved algorithm for interpreting hyperacute stroke CT brain images automatically summons an interventionist if a large vessel occlusion is detected (since large vessel occlusions may be amenable to thrombectomy).[19] A human radiologist also interprets the CT images, however, the AI typically finishes first, allowing faster access to morbidity-reducing interventional treatment. AI systems in the information specialties may soon gain sufficient competence to act autonomously in broader clinical settings, which will be highly advantageous for regions underserved by radiologists and other information specialist physicians.

AI interpretative assistants also hold certain other advantages over their human physician collaborators, such as being present and available on hospital networks 24/7 without the need for rests or breaks. Further, whereas physician thoroughness may unfortunately decline throughout the day,[78,79] AI systems perform consistently at their given levels of accuracy. In fact, if programmed to continue learning from incoming contemporaneous patients, an AI system's level of accuracy will likely increase over time.

One downside of incorporating AI interpretative assistants is that physicians may come to depend too heavily on the AI in situations where it is inappropriate. Please see Section 6 for further discussion of this possibility, and discussion of other limitations of AI's entrance into healthcare.

4.2 AI can predict important clinical outcomes from various sources

In the information specialties, AI begins with a narrow model of the patient, a model that is already contained in the investigative study. However, outside the information specialties, medically-oriented AI systems have tended to intake broader, various data sources, forming broader models of the patient (rather than working exclusively with a chest x-ray-generated model of the thorax, for example). Abstract models resulting from these broad inputs have proved useful for prediction. Similar to the pixel-by-pixel approach to medical image interpretation, AI systems in broader aspects of medicine can likewise analyze data iota-by-iota. EMRs are one excellent source of such broad data. A recent survey found that 80.5% of US hospitals were using at least a basic EMR system.[80] AI can roam freely through a patient's EMR, automatically detecting patterns and predicting things with a high degree of accuracy. For example, a 2013 study reported a deep learning system which could sift through patients' EMRs and automatically predict diagnoses of rheumatoid arthritis with moderate accuracy.[38] We will consider representative further examples of the power of AI to predict diverse clinical outcomes from transformations of broad, often routine clinical information.

A recent international collaboration used multiple machine learning strategies (including some deep learning strategies) to construct an "early warning system" for predicting mortality amongst inpatients.[81] The system extracted patterns amongst variables including certain diagnoses such as congestive heart failure and acute cerebrovascular disease, and in patients' presentation to care histories in the months prior to admission. The system was tested prospectively on a set of 11,765 patients, of whom 255 passed away. Of these, for the 69 patients (13.3%) who had been at the highest risk of passing away, the AI system accurately predicted their death 40.8 hours in advance.[81] Similar to the notion that an AI system for chest x-ray could serve to notify radiologists of the possibility of major thoracic disease,[63] the study authors of the early warning system for mortality noted that such a system could be of use to automatically notify physicians and other healthcare professionals whenever a patient exceeded a given high-risk threshold.[81]

Other work has focused on what important outcomes can be predicted from hidden patterns in routine lab data. Recent work out of the Swedish Karolinska Institutet applied various machine learning models to routine laboratory data in effort to predict outcome following traumatic brain injury.[82] Study results identified increased serum creatinine, serum glucose, and plasma osmolarity, as well as decreased serum albumin, as factors predicting a worse outcome. Other work conducted out of the University of Nottingham exposed machine learning algorithms to routine lab data for the purpose of predicting adverse cardiovascular events.[37] Compared against standard-of-care American Cardiology guidelines in a retrospective test set of 378,256 patients, the machine learning system predicted 355 more events of cardiovascular disease than did current standard of care.[37] Greater predictive accuracy will be invaluable for forecasting important outcomes such as recovery from traumatic brain injury and risk of myocardial infarction.

In addition to predicting event outcomes, routine laboratory data may also be useful for diagnosing disease. Referencing previous work demonstrating underdiagnoses of primary hyperparathyroidism,[83] Somnay et al. devised a machine learning system based on routinely available clinical data that could diagnose primary hyperparathyroidism with accuracy superior to 95%.[84] Similar to other works on early warning systems81 and alerts of possible major disease,[63] Somnay et al. suggested that their primary hyperparathyroidism-detecting system could be incorporated into EMR software to create a "best practice alert" recommending parathyroid work-up in high-risk patients.[84]

Other work has applied AI to predict response to medications. A US multicenter collaboration used a variety of machine learning approaches to make treatment recommendations for choice of antidepressant for patients with major depressive disorder, identifying a subset of patients expected to benefit from sertraline therapy relative to placebo. This benefit was observed in study results, although the sertraline-receiving patients who had been identified as optimally suited to this treatment did not experience significantly different outcomes from the other, less-optimally-suited patients receiving the same treatment.[35] Beyond antidepressants, other work employed a variety of machine learning methods to predict optimal warfarin dose, a challenging task due to its narrow therapeutic range.[36] The AI in this study had some success, and the authors concluded it could be of benefit in determining optimal dosing, especially for patients needing low maintenance doses.[36]

This research collectively demonstrates competency of various AI systems in predicting diverse clinical outcomes including mortality, adverse events, diagnosis of disease, and response to treatment. In some cases, AI prediction algorithms may be superior to existing clinical prediction guidelines, such as the widely used Framingham risk score.[37] As suggested by some, AI prediction algorithms could run alongside physicians and flash warnings when deemed relevant. In as much as competency of such systems has been demonstrated and is likely to increase, healthcare policymakers and physicians should be correspondingly enthusiastic to obtain access to AI predictive powers. AI's predictions may arise from different variables than from clinicians, and may predict different things than physicians, thus, rather than necessarily usurp physicians as lead predictors and decision-makers, AI warning systems may more likely constitute a useful and complementary second opinion. With incorporation into healthcare and exposure to high volumes of patient data, predictive machine learning algorithms will attain greater and greater accuracies. Such predictive systems may also help close the gap in care existing between areas underserved by physicians (such as Northern and remote Canada, developing countries, etc.) compared to areas with relatively abundant access to physicians. This will be a highly important outcome, since whereas information specialty tasks such as interpretation of diagnostic imaging can often be outsourced to a nearby tertiary centre, the more intimate tasks of clinical prediction, such as "which patients on my ward are at high risk of two-week mortality?" cannot be systematically outsourced in the same way. As AI systems gain increasing predictive competency, a desire for utilizing this competency should compel physicians and policymakers to consider incorporating predictive AI algorithms into healthcare.

4.3 Could physicians working alongside AI gain new insights into physiology and pathophysiology?

In receiving broad information representing a patient, such as by roaming through their EMR, AI generates a novel model of the patient that is correspondingly broader than the model contained in a diagnostic imaging study. New, AI-generated models will be interesting objects of study. In particular, models resulting from unsupervised machine learning algorithms are likely to be highly creative relative to contemporaneous human-designed models. But regardless of how an AI model is generated, physicians and policymakers will be behooved to understand its salient features. If the medical community is to sanction the adoption of a particular

AI system, it will be important for purposes of safety, potential debugging, and to understand the functioning of the system's model to ensure its robustness. As discussed above, in examining the models of AI, we may be led to new insights about the body.

AI medical systems may lead to new insights in other ways. AI's high-powered pixel-by-pixel analysis allows it see certain things that cannot be seen by physicians, thus achieving access to whole new realms of data. Within diagnostic imaging, the field of "radiomics" is emerging, which focuses on mining images for such hidden data.[85] One study demonstrated that small changes in serum potassium (as small as 0.2 mEq/L, even within the normal reference range) manifested quantifiable changes on ECG that were detectable by AI interpretation but not by human review.[86] By applying deep learning, it has been discovered that retinal fundoscopy images contain information to robustly predict a person's age, gender, blood pressure, smoking status, diabetes control, and risk of adverse cardiovascular events.[87] MR images of low-grade gliomas contain information to predict deletion of chromosomal arms 1p/19q (an important prognosticator for treatment response). Using only MR images, deep learning system acquired this knowledge with 93.3% sensitivity, 82.22% specificity, and 87.7% accuracy.[88] Perhaps most interestingly, fluorine 18 fluorodeoxyglucose PET images of the brain contain information that can be used to predict diagnosis of Alzheimer's disease 75.8 months prior to the time of eventual diagnosis with 82% specificity at 100% sensitivity.[89] Interestingly, when this model analyzed via saliency mapping to determine which features it had extracted from the data to influence its predictions, it was found that rather than relying on a specific brain location or regions that could serve as anatomic biomarkers, the system appeared to utilize data from the whole brain to inform its predictions. (The system did consider certain areas to be more influential, with some influential regions corresponding to brain regions implicated in present understandings of Alzheimer's disease.) By virtue of AI's ability to see new things, including by pixel-by-pixel analysis, and by virtue of the power of machine learning algorithms to diverge from human understandings, medical models in AI systems may lead us to new explanations of intra-body phenomena, explanations that are divergent from an otherwise evolutionary, ad hoc mode of advancement of medical knowledge.[90]

A further distinct and intriguing possibility is that AI mathematical model parameters may possibly correspond to actual biological parameters within the body. "Theory-driven" efforts in the emerging field of computational psychiatry seek to generate computation models with parameters corresponding to brain neural circuitry parameters.[91] To this end, "biophysically-realistic neural-network models"[92] have captured specific, exquisite neurotransmitter disruption caused by ketamine use;[93] detailed models of particular neuroanatomical structures such as cortico-striato-thalamic loops[94,95] have demonstrated explanatory power for various neurological and psychiatric diseases.[92,96,97] Whereas computational psychiatry approaches address the brain as an algorithm to model, this framework may map onto other organs as well. Can the pituitary gland be modelled as an algorithm for transforming serum concentrations of certain hormones, with detailed mathematical model parameters corresponding to release stimuli for the various pituitary hormones? Can a kidney be modelled as a multi-layered algorithm for filtering blood, with model parameters corresponding actions of individual nephrons, or even to arrays of particular ion pumps along individual nephrons? Indeed, a detailed computational model of nephron transportation of water and solutes was published in February, 2019.[98] Perhaps computational endocrinology and computation nephrology await us as future endpoints. For now, progress is likely to move in the direction of gradually-increasing extent of accurate modelling.

However, on the opposite end of the spectrum of model-body correspondence, it is conceivable that certain AI predictive models may have nothing to do with bodily phenomena. AI systems could hypothetically base their predictions on emergent "data" and patterns that exist only inside the model, uncorrelated to any process in or affecting the original outside system (the patient). Such predictions could even hypothetically still be robust, however, they would run the risk of being "fooled" by confounding data.[99] However, as AI predictive models achieve better and better competence, its feature maps and pattern extractions will have to be more and more consonant with the actual body itself (assuming an efficiency incentive conferred by constraints in power supply and computational storage). Thus, most likely, the parameters within the predictive models will correspond more and more closely to *some* abstract informational state about the body. Perhaps, in the extreme, something like a unified computational biology can be approached. As this field progresses, investigating model parameters to understand the salient bodily features being mapped will be of increasing importance for purposes of safety, debugging, and advancement of knowledge.

5. Limits and obstacles to physician augmentation by AI

We have explored various ways in which AI systems and physician competence will intersect. Healthcare should be keen to capture the value of AI, which will include integration of physician-AI collaborations, and, in time, perhaps attaining some reliance on autonomous AI that is uncoupled from physician oversight. However, in incorporating AI into healthcare, certain difficulties must be faced.

First, physician-AI collaboration will be limited by the rate and extent of technological advancement in medically-purposed AI systems. While some promising systems already exist, and increasing resources are being devoted to AI and specifically AI in healthcare,[10,14] much work still remains. As part of this work, we must ascertain how best to harness and implement the information provided by medical AI. (Some have noted that medicine may in general benefit from greater emphasis on its "effector arm"– i.e., knowing when and how to act on information that is available.[100]) Second, AI systems appearing competent in retrospective validation trials may require further and more extensive prospective validation on sets of real, contemporaneous patients incoming to hospital facilities. This will likely be a necessary step to achieve sufficient trust in these systems, and to foster certainty in their real-world competence. Third, and also relating to the matter of sufficient trust, is the "black box" posed by some AI systems. What happens in the deep interior layers of a neural network? What patterns are being recognized and mapped? Ability of an AI system to account for the informational, pattern-recognized basis of its output will allow greater trust in the system (if in fact the basis appears reasonable), and will help to guard against oversights via human verification at the level of pattern extraction, data mapping, etc. Ironically, it is conceivable that AI systems may be applied to help us understand other AI systems — essentially, the behaviour of a deep learning system could serve as the input information to another deep learning system. Fourth, machine learning techniques in general may fail to imbue human-level ability to reason effectively in novel circumstances, given that they will tend to be trained on the data of the past.[101] With each new patient encounter presenting a potentially novel circumstance, this limitation alone may ensure a role for human physicians in clinical medicine for the foreseeable future. Fifth, due to physician discomfort with risk and uncertainty, imperfect "assistant" AI systems could come to be inappropriately relied upon for clinical decision making. Physician dependence on AI systems must scale with competence of the given system, and AI systems validated for assistive function must not be spuriously promoted to leadership function over and above physician judgement.

Some have also suggested that current machine learning algorithms may be overhyped.[102] Relatedly, certain leading AI researchers are pursuing more advanced modes of machine learning. Neural network pioneer Geoffrey Hinton (who has been called the Einstein of AI)[102] has begun working on a new type of network called "capsule network" which may prove allow machine learning techniques of even greater power than current neural network technology.[103]

Additional concerns surrounding AI in healthcare include those of patient privacy. Machine learning algorithms generally require a vast amount of training data to refine interpretations and achieve competency. For machine learning in healthcare, the requisite data will tend to be confidential patient data, raising the question of how to expose machine learning algorithms to sufficient volumes of training data without violating patient privacy. Fortunately, sophisticated computational techniques are being devised to overcome this challenge, anonymizing patient data and providing "formal, mathematical guarantees around privacy preservation."[104] Concern has been raised that certain modern anonymization techniques, such as those used in 2018 for deidentifying information obtained from wearable devices, may be inadequate to ensure privacy.[105] Moving forward, it will be crucial to ensure privacy of any patient with medical information passing through an AI algorithm.

Lastly, there is the concern that, to whatever extent AI is relied upon in healthcare — especially without physician oversight — a computer "crash" affecting the AI would be catastrophic. Existing EMRs do occasionally crash.[106] Redundant safety back-up measures for healthcare technology must scale to be increasingly robust alongside increasing reliance on AI.

6. Limitations of this paper

This paper has focused on how AI will affect the future of medicine by intersecting with physician competence. We have not discussed the arrival of AI competence in surgical medicine, nor in pre-clinical biomedicine; unsurprisingly, the possible capabilities of AI in these fields are also evident. Robotic surgical assistants are already in widespread use worldwide.[107] Despite gains in precision and accuracy, their use generally does not significantly alter outcomes;[108] however, emerging data demonstrates potential for outcome improvement.[109] It is believed that implementing touch sensors also will help advance their effectiveness, and such developments are now forthcoming.[110] Overall, it has been proposed that "clinically feasible" autonomous surgical robots will exist before the end of this century.[111] Regarding pre-clinical biomedicine, for recent excellent detailed reviews, see Angermeuller et al. 2016[112] and Ching et al. 2018.[90]

7. Conclusion

The continued advancement of medical AI will have a tremendous and complex impact on the future of medicine and the future of human life generally. If competence of medical AI systems continues to progress, the best-performing clinical executive systems will, at some future point, almost certainly be physician-AI collaborations. Moving forward, healthcare will be increasingly confronted by the question of what to do with AI medical systems that rival physician competence. The "meeting of minds" between physicians and AI has already begun, with AI systems deployed into healthcare in numerous countries worldwide. Given the potential of AI systems to enhance accuracy and improve outcomes in manifold aspects of medicine, and to enhance the standard of care in physician-underserved regions, we may be wise from perspectives of beneficence, non-maleficence, and justice to concentrate resources on their development. Navigating the evolving, dynamic intersection of physician and AI competence will be crucial to capture benefits, mitigate risks, and achieve optimal outcomes for patients. Time will tell what will be the ultimate role of human physicians in the future of medicine.

References

[1] Stuart Russell and Peter Norvig. *Artificial Intelligence A Modern Approach Third Edition.* 2010.

[2] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig. The Microsoft 2016 Conversational Speech Recognition System. September 2016.

[3] George Saon, Gakuto Kurata, Tom Sercu, Kartik Audhkhasi, Samuel Thomas, Dimitrios Dimitriadis, Xiaodong Cui, Bhuvana Ramabhadran, Michael Picheny, Lynn-Li Lim, Bergul Roomi, and Phil Hall. English Conversational Telephone Speech Recognition by Humans and Machines. March 2017.

[4] Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming Zhou. Achieving Human Parity on Automatic Chinese to English News Translation. March 2018.

[5] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3):211–252, December 2015.

[6] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science (New York, N.Y.)*, 356(6337):508–513, May 2017.

[7] Noam Brown and Tuomas Sandholm. Libratus: The Superhuman AI for No-Limit Poker (Demonstration). Technical report.

[8] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, October 2017.

[9] Max Jaderberg, Wojciech M. Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C. Rabinowitz, Ari S. Morcos, Avraham Ruderman, Nicolas Sonnerat, Tim Green, Louise Deason, Joel Z. Leibo, David Silver, Demis Hassabis, Koray Kavukcuoglu, and Thore Graepel. Human-level performance in first-person multiplayer games with population-based deep reinforcement learning. July 2018.

[10] Yoav Shoham, Raymond Perrault, Erik Brynjolfsson, Jack Clark Openai, James Manyika, Juan Carlos Niebles, Terah Lyons, John Etchemendy, Barbara Grosz, and Zoe Bauer. Artificial intelligence index 2018 annual report. Technical report, 2018.

[11] Bernard Marr and Matt Ward. *Artificial intelligence in practice: how 50 successful companies used artificial intelligence to solve problems.* 2019.

[12] Kun-Hsing Yu, Andrew L. Beam, and Isaac S. Kohane. Artificial intelligence in healthcare. *Nature Biomedical Engineering*, 2(10):719–731, October 2018.

[13] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. Artificial intelligence in healthcare: past, present and future, 2017.

[14] Eric J. Topol. *Deep medicine: how artificial intelligence can make healthcare human again.* 2019.

[15] Jörg Goldhahn, Vanessa Rampton, and Giatgen A Spinas. Could artificial intelligence make doctors obsolete? *BMJ*, page k4563, November 2018.

[16] Ziad Obermeyer and Ezekiel J. Emanuel. Predicting the Future — Big Data, Machine Learning, and Clinical Medicine. *New England Journal of Medicine*, 375(13):1216–1219, September 2016.

[17] FDA Newsroom. FDA permits marketing of clinical decision support software for alerting providers of a potential stroke in patients — FDA, 2019.

[18] Bernard Marr. First FDA Approval For Clinical Cloud-Based Deep Learning In Healthcare. *Forbes*, 2017.

[19] Arterys FDA Clearance for Liver AI and Lung AI Lesion Spotting Software, 2019.

[20] Eric J. Topol. High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1):44–56, January 2019.

[21] Justin Petron. FDA approves stroke-detecting AI software. *Nature Biotechnology*, 36(4):290–290, apr 2018.

[22] Scott Gottlieb. Statement from FDA Commissioner Scott Gottlieb, M.D., on advancing new digital health policies to encourage innovation, bring efficiency and modernization to regulation — FDA, 2017.

[23] Arinobu Tojo. Clinical sequencing in leukemia with the assistance of artificial intelligence. *[Rinsho ketsueki] The Japanese journal of clinical hematology*, 58(10):1913–1917, 2017.

[24] Pranav Deshpande. AI could help solve the world's health care problems at scale. *VentureBeat*, 2017.

[25] Roshni Balaji. This Bengaluru startup's low-cost device is using AI to detect cervical cancer in rural Karnataka, 2019.

[26] Echo Huang. A Chinese hospital is betting big on artificial intelligence to treat patients. *Quartz*, 2018.

[27] Rita Liao. China's Infervision is helping 280 hospitals worldwide detect cancers from images. *TechCrunch.com*, 2018.

[28] Julia Powles and Hal Hodson. Google DeepMind and healthcare in an age of algorithms. *Health and Technology*, 7(4):351–367, December 2017.

[29] Thousands of Images at the Radiologist's Fingertips Seeing the Invisible, 2019.

[30] Vector Institute kicks off series of Pathfinder Projects focused on health AI adoption — Vector Institute for Artificial Intelligence, 2019.

[31] Emre Pakdemirli. Artificial intelligence in radiology: friend or foe? Where are we now and where are we heading? *Acta radiologica open*, 8(2), February 2019.

[32] Curtis P. Langlotz. Will Artificial Intelligence Replace Radiologists? *Radiology: Artificial Intelligence*, 1(3):e190058, May 2019.

[33] Maciej A. Mazurowski. Artificial Intelligence May Cause a Significant Disruption to the Radiology Workforce. *Journal of the American College of Radiology*, 16(8):1077–1082, August 2019.

[34] Emre Pakdemirli. New subspecialty or close ally emerging on the radiology horizon: Artificial Intelligence. *Acta radiologica open*, 8(3):2058460119837498, March 2019.

[35] Christian A. Webb, Madhukar H. Trivedi, Zachary D. Cohen, Daniel G. Dillon, Jay C. Fournier, Franziska Goer, Maurizio Fava, Patrick J. McGrath, Myrna Weissman, Ramin Parsey, Phil Adams, Joseph M. Trombello, Crystal Cooper, Patricia Deldin, Maria A. Oquendo, Melvin G. McInnis, Quentin Huys, Gerard Bruder, Benji T. Kurian, Manish Jha, Robert J. DeRubeis, and Diego A. Pizzagalli. Personalized prediction of antidepressant v. placebo response: evidence from the EMBARC study. *Psychological Medicine*, 49(07):1118–1127, May 2019.

[36] Zhiyuan Ma, Ping Wang, Zehui Gao, Ruobing Wang, and Koroush Khalighi. Estimate the Warfarin Dose by Ensemble of Machine Learning Algorithms. September 2018.

[37] Stephen F. Weng, Jenna Reps, Joe Kai, Jonathan M. Garibaldi, and Nadeem Qureshi. Can machine-learning improve cardiovascular risk prediction using routine clinical data? *PLOS ONE*, 12(4):e0174944, April 2017.

[38] Chen Lin, Elizabeth W. Karlson, Helena Canhao, Timothy A. Miller, Dmitriy Dligach, Pei Jun Chen, Raul Natanael Guzman Perez, Yuanyan Shen, Michael E. Weinblatt, Nancy A. Shadick, Robert M. Plenge, and Guergana K. Savova. Automatic Prediction of Rheumatoid Arthritis Disease Activity from the Electronic Medical Records. *PLoS ONE*, 8(8):e69932, August 2013.

[39] Vinod Khosla. Do We Need Doctors Or Algorithms? *TechCrunch*, 2012.

[40] Dom Galeon. For the First Time, a Robot Passed a Medical Licensing Exam, 2017.

[41] Erfan Khoram, Ang Chen, Dianjing Liu, Lei Ying, Qiqi Wang, Ming Yuan, and Zongfu Yu. Nanophotonic media for artificial neural inference. *Photonics Research*, 7(8):823, August 2019.

[42] Maricarmen Martinez and Sebastian Sequoiah-Grayson. Logic and Information. In *The Stanford Encyclopedia of Philosophy*. Edward n. edition.

[43] Angus. Stevenson. *Oxford dictionary of English*. Oxford University Press, 2011.

[44] Max Tegmark. *Life 3.0: being human in the age of artificial intelligence*. 2017.

[45] David Deutsch and Chiara Marletto. Constructor Theory of Information. Technical report, 2014.

[46] David (University of Oxford) Deutsch. *The Beginning of Infinity: Explanations that Transform the World*. Penguin Books, London, England, 2011.

[47] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. Technical report.

[48] W. Marzocchi and J. D. Zechar. Earthquake Forecasting and Earthquake Prediction: Different Approaches for Obtaining the Best Model. *Seismological Research Letters*, 82(3):442–448, May 2011.

[49] Philip F Hopkins, Lars Hernquist, Paul Martini, Thomas J Cox, Brant Robertson, Tiziana Di Matteo, and Volker Springel. A physical model for the origin of quasar lifetimes. Technical report, 2005.

[50] H M Sheffer. A Set of Five Independent Postulates for Boolean Algebras, with Application to Logical Constants. *Transactions of the American Mathematical Society*, 14:481–88, 1913.

[51] A M Turing. On computable numbers, with an application to the entscheidungsproblem. 1936.

[52] David Deutsch. Constructor Theory. Technical report, Centre for Quantum Computation, The Clarendon Laboratory, University of Oxford, and Future of Humanity Institute, University of Oxford, 2012.

[53] Richard Dawkins. *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe Without Design.*

[54] Yoshua Bengio. Deep Learning of Representations for Unsupervised and Transfer Learning. Technical report, 2012.

[55] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015.

[56] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, October 1986.

[57] Rob Garner. Celebrating 17 Years EO-1, 2017.

[58] Carles Araguz, David Llaveria, Estefany Lancheros, Elisenda Bou-Balust, Adriano Camps, Eduard Alarcon, Ignasi Lluch, Hripsime Matevosyan, Alessandro Golkar, Stefania Tonetti, Stefania Cornara, Judith Cote, Stephane Pierotti, Pedro Rodriguez, Angel Alvaro, Mateusz Sochacki, and Janusz Narkiewicz. Optimized model-based design space exploration of distributed multi-orbit multi-platform Earth observation spacecraft architectures. In *2018 IEEE Aerospace Conference*, pages 1–16. IEEE, March 2018.

[59] Dawn Chan. The AI That Has Nothing to Learn From Humans. *The Atlantic*, 2018.

[60] Saurabh Jha and Eric J. Topol. Adapting to Artificial Intelligence. *JAMA*, 316(22):2353, December 2016.

[61] Michael (Michael M.) Lewis. *The undoing project: a friendship that changed our minds.*

[62] Ramandeep Singh, Mannudeep K. Kalra, Chayanin Nitiwarangkul, John A. Patti, Fatemeh Homayounieh, Atul Padole, Pooja Rao, Preetham Putha, Victorine V. Muse, Amita Sharma, and Subba R. Digumarthy. Deep learning in chest radiography: Detection of findings and presence of change. *PLOS ONE*, 13(10):e0204155, October 2018.

[63] Eui Jin Hwang, Sunggyun Park, Kwang-Nam Jin, Jung Im Kim, So Young Choi, Jong Hyuk Lee, Jin Mo Goo, Jaehong Aum, Jae-Joon Yim, Julien G. Cohen, Gilbert R. Ferretti, and Chang Min Park. Development and Validation of a Deep Learning–Based Automated Detection Algorithm for Major Thoracic Diseases on Chest Radiographs. *JAMA Network Open*, 2(3):e191095, March 2019.

[64] Paras Lakhani and Baskaran Sundaram. Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks. *Radiology*, 284(2):574–582, August 2017.

[65] Preetham Putha, Manoj Tadepalli, Bhargava Reddy, Tarun Raj, Justy Antony Chiramal, Shalini Govil, Namita Sinha, Manjunath KS, Sundeep Reddivari, Pooja Rao, and Prashant Warier. Can Artificial Intelligence Reliably Report Chest X-Rays?: Radiologist Validation of an Algorithm trained on 1.2 Million X-Rays. July 2018.

[66] William Gale, Luke Oakden-Rayner, Gustavo Carneiro, Andrew P. Bradley, and Lyle J. Palmer. Detecting hip fractures with radiologist-level performance using deep neural networks. November 2017.

[67] Dezső Ribli, Anna Horváth, Zsuzsa Unger, Péter Pollner, and István Csabai. Detecting and classifying lesions in mammograms with Deep Learning. *Scientific Reports*, 8(1):4165, December 2018.

[68] Ali Madani, Ramy Arnaout, Mohammad Mofrad, and Rima Arnaout. Fast and accurate view classification of echocardiograms using deep learning. *npj Digital Medicine*, 1(1):6, December 2018.

[69] Yun Liu, Timo Kohlberger, Mohammad Norouzi, George E. Dahl, Jenny L. Smith, Arash Mohtashamian, Niels Olson, Lily H. Peng, Jason D. Hipp, and Martin C. Stumpe. Artificial Intelligence–Based Breast Cancer Nodal Metastasis Detection. *Archives of Pathology & Laboratory Medicine*, pages arpa.2018–0147–OA, October 2018.

[70] David F Steiner, Robert MacDonald, Yun Liu, Peter Truszkowski, Jason D Hipp, Christopher Gammage, Florence Thng, Lily Peng, and Martin C Stumpe. Impact of Deep Learning Assistance on the Histopathologic Review of Lymph Nodes for Metastatic Breast Cancer. *The American journal of surgical pathology*, 42(12):1636–1646, December 2018.

[71] Kunal Nagpal, Davis Foote, Yun Liu, Po-Hsuan, Chen, Ellery Wulczyn, Fraser Tan, Niels Olson, Jenny L. Smith, Arash Mohtashamian, James H. Wren, Greg S. Corrado, Robert MacDonald, Lily H. Peng, Mahul B. Amin, Andrew J. Evans, Ankur R. Sangoi, Craig H. Mermel, Jason D. Hipp, and Martin C. Stumpe. Development and Validation of a Deep Learning Algorithm for Improving Gleason Scoring of Prostate Cancer. November 2018.

[72] Ana Mincholé and Blanca Rodriguez. Artificial intelligence for the electrocardiogram. *Nature Medicine*, 25(1):22–23, January 2019.

[73] Awni Y. Hannun, Pranav Rajpurkar, Masoumeh Haghpanahi, Geoffrey H. Tison, Codie Bourn, Mintu P. Turakhia, and Andrew Y. Ng. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1):65–69, January 2019.

[74] Zachi I. Attia, Suraj Kapa, Francisco Lopez-Jimenez, Paul M. McKie, Dorothy J. Ladewig, Gaurav Satam, Patricia A. Pellikka, Maurice Enriquez-Sarano, Peter A. Noseworthy, Thomas M. Munger, Samuel J. Asirvatham, Christopher G. Scott, Rickey E. Carter, and Paul A. Friedman. Screening for cardiac contractile dysfunction using an artificial intelligence–enabled electrocardiogram. *Nature Medicine*, 25(1):70–74, January 2019.

[75] Hisham Daoud and Magdy Bayoumi. Efficient Epileptic Seizure Prediction based on Deep Learning. *IEEE Transactions on Biomedical Circuits and Systems*, pages 1–1, July 2019.

[76] Yogatheesan Varatharajah, Brent Berry, Jan Cimbalnik, Vaclav Kremen, Jamie Van Gompel, Matt Stead, Benjamin Brinkmann, Ravishankar Iyer, and Gregory Worrell. Integrating artificial intelligence with real-time intracranial EEG monitoring to automate interictal identification of seizure onset zones in focal epilepsy. *Journal of Neural Engineering*, 15(4):046035, August 2018.

[77] Johannes P. Koren, Johannes Herta, Franz Fürbass, Susanne Pirker, Veronika Reiner-Deitemyer, Franz Riederer, Julia Flechsenhar, Manfred Hartmann, Tilmann Kluge, and Christoph Baumgartner. Automated Long-Term EEG Review: Fast and Precise Analysis in Critical Care Patients. *Frontiers in Neurology*, 9:454, June 2018.

[78] Rebecca H. Kim, Susan C. Day, Dylan S. Small, Christopher K. Snider, Charles A. L. Rareshide, and Mitesh S. Patel. Variations in Influenza Vaccination by Clinic Appointment Time and an Active Choice Intervention in the Electronic Health Record to Increase Influenza Vaccination. *JAMA Network Open*, 1(5):e181770, September 2018.

[79] Esther Y. Hsiang, Shivan J. Mehta, Dylan S. Small, Charles A. L. Rareshide, Christopher K. Snider, Susan C. Day, and Mitesh S. Patel. Association of Primary Care Clinic Appointment Time With Clinician Ordering and Patient Completion of Breast and Colorectal Cancer Screening. *JAMA Network Open*, 2(5):e193403, May 2019.

[80] Julia Adler-Milstein, A Jay Holmgren, Peter Kralovec, Chantal Worzala, Talisha Searcy, and Vaishali Patel. Electronic health record adoption in US hospitals: the emergence of a digital "advanced use" divide. *Journal of the American Medical Informatics Association*, 24(6):1142–1148, November 2017.

[81] Chengyin Ye, Oliver Wang, Modi Liu, Le Zheng, Minjie Xia, Shiying Hao, Bo Jin, Hua Jin, Chunqing Zhu, Chao Jung Huang, Peng Gao, Gray Ellrodt, Denny Brennan, Frank Stearns, Karl G Sylvester, Eric Widen, Doff B McElhinney, and Xuefeng Ling. A Real-Time Early Warning System for Monitoring Inpatient Mortality Risk: Prospective Study Using Electronic Medical Record Data. *Journal of medical Internet research*, 21(7):e13719, July 2019.

[82] David W. Nelson, Anders Rudehill, Robert M. MacCallum, Anders Holst, Michael Wanecek, Eddie Weitzberg, and Bo-Michael Bellander. Multivariate Outcome Prediction in Traumatic Brain Injury with Focus on Laboratory Values. *Journal of Neurotrauma*, 29(17):2613–2624, November 2012.

[83] Danielle M Press, Allan E Siperstein, Eren Berber, Joyce J Shin, Rosemarie Metzger, Rosebel Monteiro, Jeff Mino, Warren Swagel, and Jamie C Mitchell. The prevalence of undiagnosed and unrecognized primary hyperparathyroidism: a population-based analysis from the electronic medical record. *Surgery*, 154(6):1232–7; discussion 1237–8, December 2013.

[84] Yash R. Somnay, Mark Craven, Kelly L. McCoy, Sally E. Carty, Tracy S. Wang, Caprice C. Greenberg, and David F. Schneider. Improving

diagnostic recognition of primary hyperparathyroidism with machine learning. *Surgery*, 161(4):1113–1121, apr 2017.

[85] Robert J. Gillies, Paul E. Kinahan, and Hedvig Hricak. Radiomics: Images Are More than Pictures, They Are Data. *Radiology*, 278(2):563–577, February 2016.

[86] John J. Dillon, Christopher V. DeSimone, Yehu Sapir, Virend K. Somers, Jennifer L. Dugan, Charles J. Bruce, Michael J. Ackerman, Samuel J. Asirvatham, Bryan L. Striemer, Jan Bukartyk, Christopher G. Scott, Kevin E. Bennet, Susan B. Mikell, Dorothy J. Ladewig, Emily J. Gilles, Amir Geva, Dan Sadot, and Paul A. Friedman. Noninvasive potassium determination using a mathematically processed ECG: Proof of concept for a novel "blood-less, blood test". *Journal of Electrocardiology*, 48(1):12–18, January 2015.

[87] Ryan Poplin, Avinash V. Varadarajan, Katy Blumer, Yun Liu, Michael V. McConnell, Greg S. Corrado, Lily Peng, and Dale R. Webster. Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nature Biomedical Engineering*, 2(3):158–164, March 2018.

[88] Zeynettin Akkus, Issa Ali, Jiří Sedlář, Jay P. Agrawal, Ian F. Parney, Caterina Giannini, and Bradley J. Erickson. Predicting Deletion of Chromosomal Arms 1p/19q in Low-Grade Gliomas from MR Images Using Machine Intelligence. *Journal of Digital Imaging*, 30(4):469–476, August 2017.

[89] Yiming Ding, Jae Ho Sohn, Michael G. Kawczynski, Hari Trivedi, Roy Harnish, Nathaniel W. Jenkins, Dmytro Lituiev, Timothy P. Copeland, Mariam S. Aboian, Carina Mari Aparici, Spencer C. Behr, Robert R. Flavell, Shih-Ying Huang, Kelly A. Zalocusky, Lorenzo Nardo, Youngho Seo, Randall A. Hawkins, Miguel Hernandez Pampaloni, Dexter Hadley, and Benjamin L. Franc. A Deep Learning Model to Predict a Diagnosis of Alzheimer Disease by Using ¹⁸ F-FDG PET of the Brain. *Radiology*, 290(2):456–464, February 2019.

[90] Travers Ching, Daniel S. Himmelstein, Brett K. Beaulieu-Jones, Alexandr A. Kalinin, Brian T. Do, Gregory P. Way, Enrico Ferrero, Paul-Michael Agapow, Michael Zietz, Michael M. Hoffman, Wei Xie, Gail L. Rosen, Benjamin J. Lengerich, Johnny Israeli, Jack Lanchantin, Stephen Woloszynek, Anne E. Carpenter, Avanti Shrikumar, Jinbo Xu, Evan M. Cofer, Christopher A. Lavender, Srinivas C. Turaga, Amr M. Alexandari, Zhiyong Lu, David J. Harris, Dave DeCaprio, Yanjun Qi, Anshul Kundaje, Yifan Peng, Laura K. Wiley, Marwin H. S. Segler, Simina M. Boca, S. Joshua Swamidass, Austin Huang, Anthony Gitter, and Casey S. Greene. Opportunities and obstacles for deep learning in biology and medicine. *Journal of The Royal Society Interface*, 15(141):20170387, apr 2018.

[91] Daniel Bennett, Steven M. Silverstein, and Yael Niv. The Two Cultures of Computational Psychiatry. *JAMA Psychiatry*, 76(6):563, June 2019.

[92] Quentin J M Huys, Tiago V Maia, and Michael J Frank. Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, 19(3):404–413, March 2016.

[93] A. Anticevic, M. Gancsos, J. D. Murray, G. Repovs, N. R. Driesen, D. J. Ennis, M. J. Niciu, P. T. Morgan, T. S. Surti, M. H. Bloch, R. Ramani, M. A. Smith, X.-J. Wang, J. H. Krystal, and P. R. Corlett. NMDA receptor function in large-scale anticorrelated neural systems with implications for cognition and schizophrenia. *Proceedings of the National Academy of Sciences*, 109(41):16720–16725, October 2012.

[94] Michael J. Frank. Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. *Journal of Cognitive Neuroscience*, 17(1):51–72, January 2005.

[95] Kevin N. Gurney, Mark D. Humphries, and Peter Redgrave. A New Framework for Cortico-Striatal Plasticity: Behavioural Theory Meets In Vitro Data at the Reinforcement-Action Interface. *PLoS Biology*, 13(1):e1002034, January 2015.

[96] Tiago V Maia and Michael J Frank. From reinforcement learning models to psychiatric and neurological disorders. *Nature Neuroscience*, 14(2):154–162, February 2011.

[97] Michael J. Frank. Linking Across Levels of Computation in Model-Based Cognitive Neuroscience. In *An Introduction to Model-Based Cognitive Neuroscience*, pages 159–177. Springer New York, New York, NY, 2015.

[98] Anita T. Layton and Harold E. Layton. A computational model of epithelial solute and water transport along a human nephron. *PLOS Computational Biology*, 15(2):e1006108, February 2019.

[99] Jiawei Su, Danilo Vasconcellos Vargas, and Sakurai Kouichi. One pixel attack for fooling deep neural networks. October 2017.

[100] Ezekiel J. Emanuel and Robert M. Wachter. Artificial Intelligence in Health Care: Will the Value Match the Hype? *JAMA*, 321(23):2281, June 2019.

[101] William W. Stead. Clinical Implications and Challenges of Artificial Intelligence and Deep Learning. *JAMA*, 320(11):1107, September 2018.

[102] James Somers. Is AI Riding a One-Trick Pony? *MIT Technology Review*, 2017.

[103] Geoffrey Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with EM routing. Technical report.

[104] Emma Bluemke. Privacy-Preserving AI in Medical Imaging: Federated Learning, Differential Privacy, and Encrypted Computation, 2019.

[105] Liangyuan Na, Cong Yang, Chi-Cheng Lo, Fangyuan Zhao, Yoshimi Fukuoka, and Anil Aswani. Feasibility of Reidentifying Individuals in Large National Physical Activity Data Sets From Which Protected Health Information Has Been Removed With Use of Machine Learning. *JAMA network open*, 1(8):e186040, December 2018.

[106] Roger Collier. Electronic medical records: preparing for the inevitable crash. *CMAJ : Canadian Medical Association journal = journal de l'Association medicale canadienne*, 186(7):493, apr 2014.

[107] Tiago Leal Ghezzi and Oly Campos Corleta. 30 Years of Robotic Surgery. *World Journal of Surgery*, 40(10):2550–2557, October 2016.

[108] Arif Ahmad, Zoha F. Ahmad, Jared D. Carleton, and Ashish Agarwala. Robotic surgery: current perceptions and the clinical evidence. *Surgical Endoscopy*, 31(1):255–263, January 2017.

[109] T. L. Edwards, K. Xue, H. C. M. Meenink, M. J. Beelen, G. J. L. Naus, M. P. Simunovic, M. Latasiewicz, A. D. Farmery, M. D. de Smet, and R. E. MacLaren. First-in-human study of the safety and viability of intraocular robotic surgery. *Nature Biomedical Engineering*, 2(9):649–656, September 2018.

[110] Chiara Bartolozzi. Neuromorphic circuits impart a sense of touch. *Science*, 360(6392):966–967, June 2018.

[111] Sandip Panesar, Yvonne Cagle, Divya Chander, Jose Morey, Juan Fernandez-Miranda, and Michel Kliot. Artificial Intelligence and the Future of Surgical Robotics. *Annals of Surgery*, 270(2):223–226, August 2019.

[112] Christof Angermueller, Tanel Pärnamaa, Leopold Parts, and Oliver Stegle. Deep learning for computational biology. *Molecular Systems Biology*, 12(7):878, July 2016.