Horizon Scanning Series

The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing

Health and Aged Care

This input paper was prepared by Frederico Girosi

Suggested Citation

Girosi, F (2018). Health and Aged Care. Input paper for the Horizon Scanning Project "The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing" on behalf of the Australian Council of Learned Academies, www.acola.org.

The views and opinions expressed in this report are those of the author and do not necessarily reflect the opinions of ACOLA.

AI and Health

Federico Girosi

Digital Health CRC

Translational Health Research Institute, Western Sydney University

Table of Contents

Introduction

Like other sectors of the economy, the health sector is composed of a large number of economic agents that continuously interact with each other. Examples of agents include patients, service providers, payers and regulatory agencies. Each interaction involves one more of agents making decisions and taking actions based on some available information. For example, a GP visit may lead to a patient initiating a treatment course that they may or may not comply with, while a payer may (or may not) pay a bill received from a provider. The actions and behaviors of the agents have consequences/outcomes that may manifest themselves after varying periods of time (outcomes of surgery may be known quickly, but outcomes of changes in lifestyle may take much longer). In addition, decisions and behaviors are rarely part of a "one-shot game": they tend to be part of a chain of related events, with outcomes depending on the whole sequence rather than the individual event (for example missing one dose of medication may be irrelevant, but missing a whole lot is not).

The environment in which decisions are made and behaviors take place has the following characteristics:

- there is considerable amount of heterogeneity among all agents involved;
- decisions are often made under large amounts of uncertainty;
- decisions are complex, and the optimality of choices is bounded by limits of human cognitive abilities as well as limits on the time allowed to reach a decision;
- decisions must be made taking in account constraints of finite resources.
- behaviors of patients, providers and institutions are resistant to change, and effort may be needed to sustain change.
- the modality of transmission of information to and from agents is often as important as the content.

Given these characteristics of the health care system it follows that improvements can be realized in at least two different ways:

1. improving the decision-making process. This includes removing some of the uncertainty, so that decisions are more grounded in knowledge and evidence, and improving the optimality of the choices, so that humans are not bounded by cognitive limits anymore;

2. helping humans to change behavior, if they are so inclined, and to sustain the change.

Independently of how exactly we define AI, it can certainly play a role in both the areas above. By analyzing large amount of data and extracting meaningful patterns AI can remove some of the uncertainty that is so pervasive in health care [1], leading to agents making better informed decisions and to an improvement of overall efficiency [2]. In addition, by utilizing human-like interfaces and means of communication AI can also play a significant role in helping patients and providers to modify their behaviors in sustainable ways. The specifics of those roles are described in the following section.

The Role of AI in Health

The possibilities to apply AI in health are numerous, and an excellent overview can be found in [3]. Readers interested in looking at the latest applications are recommended to look at the proceedings of the annual conference Artificial Intelligence in Medicine [4]. In the following section we cover some of the most promising areas, where much progress is expected in the short to medium term.

The Personalization of Medicine

Populations are enormously heterogenous in regard to health risks and behaviors. Every person's health and wellness trajectory is different, influenced by biological, environmental and social factors. Many of us will suffer from chronic diseases and some of us will experience adverse events as part of the care we receive. Among those who suffered a transport or work accident the timing and degree of the recovery will be different from one another. Providers also differ greatly in the quality of care they provide and in their propensity to make errors and possibly over or under service.
A health care system that is unable to take this heterogeneity in account ends up delivering care which is generic, rather than personalized. This leads to some patients receiving care they do not need (costing more and possibly creating harm) or that is not timely or appropriate (creating harm and possibly costing more in the longer term). AI can deliver knowledge-based solutions tailored to a person's circumstances, history and preferences that can reduce these inefficiencies and help to provide the right care at the right time.

Prevention

Great potential is found in the area of prevention. Since risk tend to be highly concentrated in small population sub-groups there is a clear role for AI to identify the relatively few individuals at high risk of developing a chronic condition, experiencing an adverse event or recovering poorly from an injury, or to identify the providers more likely to be associated with care of low value or quality. Combining individual health data with expert opinion AI solutions can answers questions of the following type: “What is the risk that person X will experience (or have characteristic) Y?” In this context Y could be both negative events (developing a chronic condition, being re-admitted to hospital, not recovering well from accident or trauma, requiring extraordinary amounts of resources, being a fraudulent provider) and positive events (benefiting from a treatment or an intervention, being a good candidate for a randomized control trial, being a provider of high quality).
AI systems of this type share with humans the ability to learn from examples. Usually they are presented with large numbers of instances of a problem and are able to “generalize”, that is to extract complex rules underlying the dynamic of the system. To the extent that these rules can be simplified and made understandable to humans the AI solutions will not only predict certain events, but also help to explain why they happen (for example which combination of past events and variables may contribute to it).

Diagnosis, Prognosis and Treatment

At the point of care much of the services provided to patients could be personalized using intelligent and knowledge-based systems [5] [6]. At the moment providers are limited in their ability to diagnose not only by time and resources (there is only so many tests one can run) but also by the fact that medical and scientific knowledge increases at such a high pace [7] that it has become unpractical to remain up to date with current knowledge. AI systems can perform critical tasks at the point of care:

1. They can process massive amount of medical literature and guidelines using Natural Language Processing and provide clinicians with summaries of evidence and recommendations which are relevant to a specific patient;
2. They can sift through large amounts of patients’ records and find cases “that looks like this”. This is particularly useful in case of rare conditions and can help not only the diagnosis but also the prognosis, giving patients a better picture of how their future trajectories may look like.
3. They can help to interpret the results of diagnostic tests. Imaging and Pathology are two fields that are already changing thanks to the deployment of AI systems and machine learning technologies [8], where already in some cases machines outperform humans [9]. Much potential lies

in relying on automated system for the performance of simpler tasks (such as detecting lung nodules or counting cells), freeing radiologists and pathologists to perform complex tasks, that requires the type of human reasoning AI system are still unable to perform.

4. They can help to design more personalized treatment plans and to provide treatment recommendations, based on the knowledge extracted from the analysis of large numbers of patients. In addition, AI solutions can help patients to adhere to a treatment regime by engaging with the patient through virtual assistants and intelligent reminder systems.

Value, Quality and Safety

Not all Australians receive safe, effective, or high value care. It has been argued that in many respects we are "flying blind" in terms of our understanding of the value of the care delivered, the performance of organizations and teams, and the optimality of resource allocation [1].

The issue of value, which is of particular importance to payers, has been traditionally dealt within the framework of health economics or by combining expert opinions with analysis of data [10] [11]. The promise of AI is to provide a much more nuanced vision of value, that takes in account the heterogeneity of outcomes and recognizes that value means different things to different individuals.

A key component for the analysis of value of care is the measurement of health outcomes, which is notoriously difficult to do based on administrative data, it may require the administration of surveys and it is still hard to do even with access to patient records. AI solutions have the potential to communicate with patients via chatbots and to extract valid PROMs (Patient Reported Outcome Measures) from analysis of text gathered via SMS, e-mails or social media. Such measures could be not only used to study value and outcomes of interventions but could also be used to provide real-time feedback to the clinicians treating those patients, feeding into the next generation of Clinical Decision Support Systems (CDSS).

CDSS are systems that combine existing medical knowledge and patient observations, often integrated with electronic health records, that can assist the clinical decision-making process at the point of care. Current systems have the potential to increase clinician adherence to guidelines, facilitate communication between providers and patients, improve the quality and safety of medication prescribing, and decrease the rate of prescription errors [12] [13]. The next generation of CDSS will benefit from AI in several ways:

1. the same way medicine is becoming personalized for patients, the practice of medicine can become more personalized for clinicians. At the moment CDSS do not take in account clinicians' preferences, may not integrate well

with their workflow and are not adaptive. AI can introduce these features, that tend to be barriers to adoption of CDSS, and become an enabling factor for improving clinical practice and reducing clinical variation.

2. AI enabled systems can collect information about the patient, summarize it and feed it back to clinicians in a continuous loop. This is currently missing from most clinical practices, that have limited abilities to evaluate the health status of their patients and the quality of care provided.
3. AI enabled systems could also allow clinicians to compare themselves to others, making "fair comparisons" that take in account the varying composition of the mix of patients they treat.
4. AI enabled systems in acute or rehab care can also monitor patients and predict in advance which ones might be about to experience an adverse event or are deviating significantly from their expected recovery path.

Consumer Empowerment

In the words of Leonard Kish, the blockbuster drug of the century is the engaged patient [14]. As the population continues to age and disease becomes increasingly chronic rather than acute, the key question is how we support consumers with different characteristics and personalities to engage in the management of their own health and wellness and stay healthy and independent for as long as possible.

Designing programs to remain healthy as well as to manage current health conditions requires individuals not only to change current behaviors (such as dieting and smoking), but also to adopt completely new ones (such as managing new, complex medications). These tasks are notoriously difficult, and that is where AI systems will at the same time meet their greatest challenge and have the potential to achieve high impact.

A key role played by AI will consists precisely in the engagement of the consumer, by choosing and adapting the way they interact with them, recognizing moods and needs, and continuously shifting the interaction to maintain a high level of engagement. Next is the provision of information: AI systems have the ability of provide context-specific information, that depends on the individuals, where they live, their preferences and their health history. Here it may lie their greatest contribution: unlike static, not-adaptive systems, intelligent chatbots and virtual assistants may be able to provide actionable and timely information, answering questions such as "May I have ice cream now?" rather than reciting the amount of calories intake one is expected to have daily.

The range of applicability of consumer empowering applications is wide and goes from the personalized management of chronic conditions such as diabetes or heart disease, to the management of medicines, for individuals on polypharmacy, to applications that support behavioral health care in areas related to mental health [15]. In this area AI solutions have much potential to deliver Intelligent Virtual Agents [16], computer-controlled characters that can interact with humans, are capable of affective interaction and can recognize human emotions [15].

Business Efficiency and Resource Allocation

There is great potential for AI to improve the efficiency of businesses by an intelligent replacement or enhancement of specific tasks. Consider tasks such taking clinical notes, coding patient records according to coding systems such as ICD or SNOMED or handing over clinical information from one heath practitioner to another [17]. These tasks are time consuming [18] and critical for the provision of good quality care [19] as well as for appropriate billing. Far from replacing humans, AI solutions involving voice recognition and Natural Language Processing capabilities can summarize conversations, convert information from unstructured (say text) to structured format (say tables, codes and check boxes) as well as provide intelligent recommendations (for example: "records that look like *this* tend to be coded as *that*"). Systems of this type would allow humans to spend more time in productive activities, enhancing both the productivity of the workplace as well as the quality of care, at the same time providing an element of consistency and reducing random variation.

Other opportunities for AI to increase business efficiency lie in the areas of fraud and error detection [20] as well as smart claiming management. This area is of interest to both providers and payers, who are locked in a constant struggle: providers worry about under-billing and payers worry about being over-billed. This usually entails a large proportion of claims being labeled "unusual" and being reviewed, often by auditors and by hand, with limited guarantee to find an actual error [21]. While analytics and intelligent solutions started to appear [22], AI can provide a variety of solutions that combine human intelligence (for example using insights about human behavior) with the analysis of massive administrative data. These systems would contribute to create a more transparent and efficient environment that catches errors before they are committed (for example by use of alert systems on the provider side), where only few records need to be manually reviewed, and where all parties are in agreement on the adjustment to be made.

Recent Progress

One of the areas of AI that has seen most progress is the processing of unstructured data, which includes text as well as images and speech. Unstructured data is pervasive in health, and it is commonly agreed that it constitutes the majority of health data. The ability of a computer program to summarize, categorize and interpret unstructured data highlights the "I" in AI, since these are tasks at which the human brain excels [23]. AI solutions that use these types of data are able generalize from examples like humans do ("*this* condition usually looks like *this*"), and this capability has allowed to build systems with performance similar to humans, or with performances good enough that they can support clinical decision-making.

The most common sources of unstructured data are clinical notes/reports and images. Advances in both image processing and Natural Language Processing [24]are already transforming the field of radiology and imaging [25][25] [9] [8] (see [26] for a review of applications of AI to radiology), with important implications for patients with heart disease [26] [27] [28, 29], cancer [30] and conditions affecting the eye [31]. An impressive example is Arterys[1], a cloud-based AI radiology assistant that has obtained FDA approval in the US. Not only Arterys can reach at least the same accuracy of humans, but it takes only 15 seconds to reach a conclusion when a human would take 30 minutes to an hour [32].

More generally, AI solutions can use a variety of data of different types, both unstructured and structured (like clinical variables, pathology test results and genetic information) to assist doctors in the diagnosis process of different disease [2] [6]. For example, in oncology the literature on the accuracy of these systems is favorable, although stronger evidence is still needed [33] before they can enter mainstream clinical practice [34].

Much of the progress in the applications of AI to health derives from the sensational progress experienced by the field of machine learning in the past few years, with the advent of deep-learning. Deep-learning methodologies have been applied to a large variety of domains, including health. An obvious application of deep-learning and neural-network models is to prevention, where people ask questions like "who is at risk for developing/experiencing X?", where X is usually a negative event. For example, machines proved to perform better than humans in predicting suicide [35] and in discriminating malignant breast lesions [36], and in predicting who is at risk for a variety of chronic conditions [15, 37].

[1] https://arterys.com/

A similar problem is the one of prognosis, where one asks questions of the following type: "given a diagnosis, how does the trajectory of the patient look like? what is the survival rate for a patient that looks like this? what is the probability of recurrence (for cancer)?". Often a combination of clinical and genetic information it is used (see [30] for a review), and while it is true that there has been little penetration of these techniques in clinical practice the progress in this area has been remarkable and it is expected only to improve. This is best exemplified by the fact that large, innovative companies such as IBM and Google have heavily invested in health. Systems such as IBM Watson [38] hold enormous potential in harnessing unstructured data, such as clinical notes, and using them to provide treatment recommendations of cancer, although they are not necessarily near to become mainstream [39]. Some solutions by Google DeepMind [40] are being used in hospital care, in the UK. DeepMind can pull together patient data and present them to nurses in a ward, saving considerable time, and send alert of impending patient deterioration. While the version currently in use with the UK NHS does not contain much AI at the moment, the fact that it has been adopted by a major stakeholder constitutes enormous progress for the field.

Progress has also been made, although not nearly as extensively, in the area of spoken and text-based dialogue systems (DS) [41]. DSs are of particular importance in the health sector, where the quality of interaction between human and virtual agents is crucial to maintain engagement. Substantial innovation has taken place regarding applications of DS to Mental Health [41], which range from the design of virtual affective agents, that can help patients with depression or autism [16] to the delivery of interventions [41].

DS are often embedded in mobile smart apps for wellness and personal health management, an area that has seen an explosion of activities over the past few years. However, how much of that constitutes actual progress is debatable. To begin, not every mobile health app necessarily qualifies as application of AI, since many apps do not really have an "intelligent" component (such as the ability to adapt to the user, to learn from past behaviors or data sets, to interact as a human or to perform some form of reasoning). In addition, the evidence that these apps provide any health benefits is often scant [42]. Progress has been made, however, on the measurement of the quality of apps (such as the MARS scale in Australia [43]) or on their certification. In the latter area, the US has led the way by allowing the FDA (Food and Drug Administration) to approve mobile apps the same way they approve drugs (hence the name "prescription apps") [44] [45], and the NHS in the UK has a library of "trusted" digital tools, some which may have an AI component.

The range of applications in this area is wide (see [45] for an extensive review). Many solutions are devoted to help consumers to change behaviors and better manage chronic conditions, in particular diabetes and heart disease. Apps often provide support to change dietary habits, to stop smoking/drinking and to increase levels of physical activity. Many solutions include an element of monitoring, looking for exacerbation of symptoms, and may include interaction with health providers, such as nurses, dieticians and mental health specialists [46]. An interesting opportunity offered by mobile apps is that consumers may be more prone to disclose information to an app rather than a human [47], opening the way for better informed services.

To repeat, many of these apps may have only a small element of AI, and often are mostly passive devices, but it is clear that they are becoming increasingly intelligent and able to adapt to the needs of the consumers and interact in a more human form.

There are clearly other areas in which progress has been made, one of them being the detection of fraud and waste in health insurance. There is a consensus that large amounts of waste, fraud and error take place in the health care system systems [21] and this is matched by a growing literature in this area [48] [49] [50] [51] [52] [53] [54] [55] [56, 57]. Much of the work is based on the detection of outliers and "unusual behaviors", since often it is not known in advance how fraud may look like. Some of the work also relies on recent in advances in graph and social network analyses, that allow to extract from large data sets suspicious and unusual relationship among different agents in the health care system. Some systems rely on the analysis of large amount of data, and "discover" previously unknown types of fraud, while others use sets of rules that have been designed by experts in the field and that mimic human intelligence. One example of this type of application is HIBIS, a product developed by the Australian company Lorica Health, that has found wide applicability in Australia health insurance market.

Gaps Analysis

Regulatory Barriers, Data Access and Infrastructure

By far the greatest challenge faced by researchers in AI and Health is access to individual level data, on both the patient and the provider side. This has consequences on the development of tools and solutions as well on the training of data scientists, as outlined in the next section, and if not addressed it will greatly hamper the applications of AI to health. While the fragmentation of the health system, that leads to health data being scattered and siloed across different data sets and institutions (more in Australia than in New Zealand), is a contributing factor to lack of access, the main culprit is the regulatory and legal environment. This aspect has been well documented in the literature [58] [59] [60] and eloquently articulated in the Productivity Commission Inquiry "Data Availability and Use", especially in the section appropriately named "What holds us back?" [60].

It is important to notice that while it is true that legislation around privacy and data use is often outdated and restrictive in modern terms, it is often the regulations and guidelines set in place by specific bodies that limit access the data. As suggested by the Productivity Commission, it is then possible that researchers may not be able to take advantage of what is legally permissible because of a culture of risk aversion in the public sector, mixed with misinterpretation of the law and an incorrect reading of the attitude of the general population, which is likely to be more open to data sharing than what regulatory agencies may perceive [61] [62] [60].

A contributing factor to this restrictive data environment is lack of specificity of regulations and guidelines. A prime example is the National Statement on Ethical Conduct in Human Research [63], the document that sets national standards for use by any individual or institution conducting human research and that is constantly used by Ethics Committees to aid their decisions and by researchers to make sure their research protocols stand on solid ethical ground. The National Statement is a truly inspiring and well thought document that lays ethical *principles* with great clarity. However, it is not meant to be an operational document, and leaves ethics committees and researchers alike with no practical advice around data access. For example, industry and government bodies may want to share and link de-identified data with researchers, but there is no guidance on what a de-identification protocol may look like, on what is the risk of re-identification under different circumstances, and what is an acceptable level of risk. As a result, data sharing, rather than being restricted just enough to preserve privacy, may not take place at all.

While the picture painted here and in much greater detail in the excellent Productivity Commission Inquiry report[4] may appear negative, the upside is that the gaps in data access could be closed relatively quickly given the collective will to do so, even without major legislative changes. The technology to link data, maintain it secure, guarantee privacy and provide access to researchers is largely available, and Australia has invested heavily in infrastructure programs such as the Population Health Research Network (PHRN) and in data linkage agencies across most states and territories. Costs of access and linkage remain a barrier [64], but not a structural one, and one that would be naturally lowered if the entire process of getting data access were simplified and inefficient replication of ethics applications and data repositories could be mitigated.

What is required then to bridge these gaps is a massive collaborative effort among ethics committee, government agencies, and teams of ethicists and lawyers, aimed to design simplified processes and clear guidelines for data sharing.

Skills and Training

Even a cursory look at the courses offered by Australian and New Zealand Universities shows that there is no shortage of training opportunities in the areas of AI and Data Science, which are increasing over time. This alone, however, it Is not sufficient to guarantee that these countries are ready to educate the next generations of researchers and developers that can apply AI and data science methods to health. Several obstacles stand in the way. One is the lack of easy access to health data, described in the earlier section, that makes it difficult for higher education students to get crucial practice and develop skills specific to the health area. PhD students at well-funded institutions, with established research programs in areas such a linked data, may be able to take advantage of existing projects and funding to gain access to valuable data, but this type of training remains inaccessible to most students interested in the field. As a consequence, we may end up with motivated students moving to different applications of AI than to health or graduating without adequate training. It is also possible for students to gain experience using data sets from other countries, such as the US, which can be easily downloaded for free or for modest fees[2]. This leads to an enormous opportunity cost, since society would have been better off if these students trained and published off Australian or New Zealand data sets.

Another obstacle to the training of the next generation of AI scientists working in health is the potential lack of domain knowledge. Developing applications of AI in health requires a keen understanding of human health, population health, human behaviors and regulatory environment. With very few exceptions[3], higher education in AI and data science tends to be generic, with very little intersection with health. Closing this gap may be possible not only by developing degree programs which are more specific to heath, but also by facilitating interactions between students and industry sectors, for example taking advantage of institutions such as CSIRO, programs such as the Australian Cooperative Research Centers, or initiatives such as those supported by Callaghan Innovation, New Zealand's innovation agency.

---

[2] One year of data from the US HCUP National Inpatient Sample, containing about 7 million hospitalizations, costs about $100 to students.

[3] For example, UNSW has recently started offering, in addition to their Master in Health Data Science, a suite of Health Data Science professional development courses (https://cbdrh.med.unsw.edu.au/professional-development).

The Next 10 Years

The future of evolution of any field has drivers, factors that push it forward, and obstacles that hold it back. In the case of AI and health, a key driver is the ever-increasing proportion of GDP that is spent in health, driven by a combination of technological [65]and demographic change [66] which raises serious sustainability issues. This has created a sense of urgency for finding ways to contain costs. At the same time the proportion of people with multiple chronic conditions who take multiple drugs has been rapidly increasing, demanding more care and of better quality. Given the great potential of AI to improve both efficiency and quality of care one can only expect to see a great demand for AI-based solutions. Both Australian and New Zealand government are certainly investing in the sectors, and markets and investors seem to confirm these expectations, since the market for AI application in health is expected to experience very fast growth [67].

While this is true for most countries, levels of readiness for change and maturity vary. In a recent survey of business leaders assessing AI maturity Australia ranked last among 7 countries (highest were China, India and Germany, followed by US, France, UK and Australia) [68]. This factor, together with the gaps and barriers described in the previous section, will be clearly holding back progress in the region.

Within this complex landscape one can expect two important types of activities going forward and in parallel. On one side, there is a host of AI solutions that have been developed already and are suitable to be trialed, refined and deployed, and huge opportunities lie in taking advantage of those. On the other side there is highly fruitful areas that still need serious development and may require large shift in model of cares to be implemented. We describe them briefly below.

Trialing and Deploying

Solutions that analyze images or clinical notes and help clinicians to diagnose a condition [69, 32], or to present the patient with a personalized prognosis, or that help to design a personalized treatment plan, are quite in mature state and we would expect much progress in terms of adoption along technologies of this type [25][25][9][8][26] [26] [27] [28, 29] [30] [31]. Solutions of this type tend to be quite specific and aim to assist clinicians in either performing the same task in less time or performing it more accurately. Therefore, while they still require some change in practice style, they are expected to be seen as less threatening and more likely to be adopted.

Solutions aiming to empower consumers, deliver mostly by mobile apps, abound and are in different state of readiness to be trialed, going from the FDA approved BlueStar[4] app for diabetes management to a host of unrated apps. The public has certainly an appetite for intelligent solutions to manage their health and wellness, but the proliferation of apps and lack of certification and quality measurement is a hindering factor at this point. Therefore, we would expect to see the deployment and trial not of single apps, but of one-stop-shop apps, or “eco-systems of apps”, that allow consumers to manage all aspect of their health and wellness in one place.

Within this universe we would also expect to see applications that do not require interaction with clinicians or medical records to leap forward, since they will find a friendlier regulatory environment. In particular, wellness and preventative solutions that help consumers to adhere to a physical activity regime, or that assist the monitoring of diet by providing actionable and personalized information, seem very well placed.

Developing

The large investments in electronic medical record systems and the recent advances in interoperability standard such as FHIR [70] [71], the likely shift to value-based payment systems that reward outcomes, and the desire of clinicians and nurses to spend less time in administrative tasks and more time in providing high quality care suggest that medical practice is likely to change significantly over the next decade, and AI solutions are a key enabler of this change.

A key contribution of AI will be the delivery of the next generation of clinical decision support system (CDSS). Intelligent CDSS will be able to capture the characteristics of the patient as well as the style of practice of the provider. To begin with, current advances in NLP and speech recognition will allow devices to listen to the conversation between patient and provider, capture both clinical and billing codes, update the patient’s records and summarize it: providers will only have to review it (and possibly modify it) and will then be able to fully dedicate themselves to their patients.

Smart apps will help consumers to collect patient reported outcome measures (PROM) without having to fill in questionnaires, using voice interaction or perhaps monitoring a person activities or interaction with others. Critically, these measures will be fed back to the providers, who will be able to obtain not only a clear

[4] https://www.welldoc.com/product/

picture of how their patients are doing but also whether they are following a trajectory that is within the norm for patients with those characteristics.

The ability to use predictive models to compare the observed status of the patient with the expected status is a key contribution of AI to health. Applied in an acute setting it will allow providers to prevent the deterioration of patient or capture an impeding negative event. Applied in primary care or in a rehabilitation setting it will allow to capture individuals who are not on an optimal health trajectory and intervene before it is too late.

In addition, the paradigm of measuring quality by comparing the observed with the expected is a powerful one, because it relies on predictive models whose accuracy can be objectively measured, and therefore leave less space to ambiguity and interpretation. Progress in this area it will clearly be slowed down by the barriers described in the previous section, but at the same time we expect some of those barriers to be lowered in the next decade.

Where those barriers are already lower, we should expect to see faster progress. This is the case, for example, of applications of AI to business efficiency. The technologies needed to intelligently assist medical coders and to catch potential billing errors before they take place, or to detect errors and fraud in real time, are certainly available. Businesses willing to experiment with AI solutions in this space have access to their own historical data and do not face difficult clinical governance and ethical questions. Therefore, given the intense cost pressure on both payer and provider side we would expect significant progress along these lines over the next decade.

Resources and Call to Action

- **Government**: while government investments in this area are useful, leadership and coordination would be even more useful. Many of the barriers described in the section above cannot be overcome unless government takes a proactive role. Listening and acting on the recommendation of the Australian Productivity Commission would be a perfect starting point. Making innovation part of the culture of government bureaucracy is a challenging but necessary step, and the same can be said of the need of better communication and collaboration among different components of government, from federal to jurisdictional level. It is also crucial for governments to develop reliable mechanisms for listening to the will of the people: culture is changing and people's attitudes regarding research and privacy may be different of what the government may currently think.

- **Industry**: industry has clearly showed much interest in the applications of AI in health. While it is true that the current regulatory environment makes it difficult for industries to develop, trial, and deploy solutions, progress is also hindered by lack of readiness on the industry side. Managers will need to become much more knowledgeable with AI and analytics in order to reap the benefits that AI may bring to them.

- **Education**: researchers working on AI solutions for health need a very solid domain knowledge of health, that ranges from the clinical to the system level view. At the moment there are not many interdisciplinary programs that provide this type of training, so they will have to be developed in order to make progress. In addition, in order for the health workforce to take advantage of AI solutions they will also need access to training, that needs to be developed.

- **Individuals**: consumers and providers alike are not passive users of AI solutions. Consumers can play a crucial role in speeding up the adoption of AI in health by becoming more informed about the benefits and implications of the personalization of medicine and by making their views known. There will have to be a realization that in order to achieve any level of personalized medicine some personal data will have to be analyzed, and an honest and informed discussion between consumers and policy makers regarding the use of data will have to take place.

References

[1] U. Srinivasan, S. Rao, D. Ramachandran and D. Jonas, "Flying Blind: Australian Consumers and Digital Health, Volume 1.," Australian Health Data Series, Health Market Quality Research Program., Sydney, 2017.

[2] E. Loh, "Medicine and the rise of the robots: a qualitative review of recent advances of artificial intelligence in health," *BMJ Leader,* 2018.

[3] JASON, "Artificial Intelligence for Health and Health Care," JSR-17-Task-002, 2017.

[4] A. t. Teije, C. Popow, J. Holmes and L. Sacchi, Proceedings of the 16th Conference on Artificial Intelligence in Medicine, June 21-24, Vienna, 2017.

[5] N. V. Chawla and Darcy A. Davis, "Bringing Big Data to Personalized Healthcare: A Patient-Centered Framework," *Journal of General Internal Medicine,* vol. 28, no. 3, 2013.

[6] J. Fei, "Artificial intelligence in healthcare: past, present and future," *Stroke and Vascular Neurology,* vol. 2, no. 4, 2017.

[7] L. Bornmann and R. Mutz, "Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references," arXiv:1402.4578.

[8] S. Jha and E. J. Topol, "Adapting to Artificial Intelligence," *JAMA,* vol. 316, no. 22, 2016.

[9] K. Yu, "Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features," *Nature Communications,* vol. 7, no. 12474, 2016.

[10] J. Schwartz, B. Landon, A. Elshaug, M. Chernew and J. McWilliams, "Measuring low-value care in Medicare," *JAMA Internal Medicine,* vol. 174, no. 7, 2014.

[11] A. Elshaug, A. Watt, L. Mundy and C. Willis, "Over 150 potentially low-value health care practices: an Australian study," *Med J Aust,* vol. 197, no. 10, 2012.

[12] L. Moja, K. Kwag and T. Lytras, "Effectiveness of Computerized Decision Support Systems Linked to Electronic Health Records: A Systematic Review and Meta-Analysis.," *American Journal of Public Health,* vol. 104, no. 12, 2014.

[13] T. Bright and others, "Effect of Clinical Decision-Support Systems," *Annals of Internal Medicine,* vol. 157, no. 1, 2012.

[14] L. Kish, "The Blockbuster Drug of the Century: An Engaged Patient," [Online]. Available: http://healthstandards.com/blog/2012/08/28/drug-of-the-century/.

[15] G. Bedi, F. Carrillo, G. Cecchi and others, "Automated analysis of free speech predicts psychosis onset in high-risk youths," *NPJ Schizophr. ,* vol. 1, 2015.

[16] D. D. Luxton, Artificial intelligence in behavioral and mental health care, Academic Press, 2015.

[17] H. Suominen, "Capturing patient information at nursing shift changes: methodological evaluation of speech recognition and information extraction,," *Journal of the American Medical Informatics Association,* vol. 22, no. e1, p. e48–e66, 2015.

[18] G. . Hripcsak, D. K. Vawdrey, M. R. Fred and S. . Bostwick, "Use of electronic clinical documentation: time spent and team interactions," *Journal of the American Medical Informatics Association,* vol. 18, no. 2, pp. 112-117, 2011.

[19] R. Amarasingham, L. Plantinga, M. Diener-West, G. DJ and N. Powe, "Clinical Information Technologies and Inpatient OutcomesA Multiple Hospital Study.," *Arch Intern Med.,* vol. 169, no. 2, 2009.

[20] K. Flynn, "Financial fraud in the private health insurance sector in Australia: perspectives from the industry," *Journal of Financial Crime,* vol. 23, no. 1, 2016.

[21] S. Hehner, "Artificial Intelligence in Health Insurance," McKinsey & Company, 2017.

[22] U. Srinivasan and B. Arunasalam, "Leveraging big data analytics to reduce healthcare costs," *IT professional,* vol. 15, no. 6, 2013.

[23] J. Hirschberg and C. Manning, "Advances in natural language processing," *Science,* vol. 349, no. 6245, 2015.

[24] C. Friedman and N. Elhadad, "Natural Language Processing in Health Care and Biomedicine," in *Biomedical Informatics*, E. Shortliffe and J. Cimino, Eds., Springer, London, 2014.

[25] R. Gillies, P. Kinahan and H. Hricak, "Radiomics: Images Are More than Pictures, They Are Data," *Radiology,* vol. 278, no. 2, 2016.

[26] K. Johnson and others, "Artificial Intelligence in Cardiology," *Journal of the American College of Cardiology,* vol. 71, no. 23, 2018.

[27] S. Massalha and others, "Decision Support Tools, Systems, and Artificial Intelligence in Cardiac Imaging," *Canadian Journal of Cardiology ,* vol. 34, no. 7.

[28] S. Dilsizian and E. Siegel, " Artificial intelligence in medicine and cardiac imaging: harnessing big data and advanced computing to provide personalized medical diagnosis and treatment.," *Current cardiology reports,* vol. 16, no. 1, 2014.

[29] I. Paschalidis, "How machine learning is helping us predict heart disease and diabetes," *Harvard Business Review,* 2017.

[30] K. Kourou and others, "Machine learning applications in cancer prognosis and prediction," *Computational and structural biotechnology journal,* vol. 13, 2015.

[31] V. Gulshan and others, "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs," *JAMA,* vol. 316, no. 22, 2016.

[32] B. Marr, "First FDA approval for clinical Cloud-Based Deep Learning in Healthcare," 2017. [Online]. Available: https://www.forbes.com/sites/bernardmarr/2017/01/20/first-fda-approval-for-clinical-cloud-based-deep-learning-in-healthcare/#459ed67c161c. [Accessed 24 7 2018].

[33] P. Lisboa and A. Taktak, "The use of artificial neural networks in decision support in cancer: a systematic revie**w,"** ***Neural Networks,*** **vol. 19, no. 4, 2006.**

[34] J. Chen and S. Asch, "Machine learning and prediction in medicine—beyond the peak of inflated expectations," *The New England Journal of Medicine 376.26,* vol. 376, no. 26, 2017.

[35] C. Walsh, J. Ribeiro and J. Franklin, “Predicting risk of suicide attempts over time through machine learning,” *Clinical Psychological Science,* vol. 5, no. 3, 2017.

[36] T. Ayer and others, “Breast cancer risk estimation with artificial neural networks revisited: discrimination and calibration,” *Cancer 116.14 (2010): 3310-3321.,* vol. 116, no. 14, 2010.

[37] S. Griffiths, “This AI software can tell if you're at risk from cancer before symptoms appear,” [Online]. Available: https://www.wired.co.uk/article/cancer-risk-ai-mammograms. [Accessed 23 7 2018].

[38] D. Arnold and T. Wilson, “What doctor? Why AI and robotics will define new health,” PWC, 2017.

[39] S. Somashekhar and others, “Double blinded validation study to assess performance of IBM artificial intelligence platform, Watson for oncology in comparison with Manipal multidisciplinary tumour board–First study of 638 breast cancer cases,” *Cancer Res* , Vols. S6-07, 2017.

[40] J. Powles and H. Hodson, “Google DeepMind and healthcare in an age of algorithms,” *Health and Technology* , vol. 7, no. 4, 2017.

[41] S. Hoermann, K. McCabe, D. Milne and R. C. RA, “Application of Synchronous Text-Based Dialogue Systems in Mental Health Interventions: Systematic Review,” *Hoermann S, McCabe KL, Milne DN, Calvo RA. Application of Synchronous Text-Based Dialogue Systems in Mental Journal of Medical Internet Research,* vol. 18, no. 8, 2017.

[42] O. Byambasuren and others, “Prescribable mHealth apps identified from an overview of systematic reviews,” *npj Digital Medicine,* vol. 1, no. 1, 2018.

[43] S. Stoyanov and others, “Mobile app rating scale: a new tool for assessing the quality of health mobile apps,” *JMIR mHealth and uHealth,* vol. 31, 2015.

[44] N. Bilbrough, “The FDA, congress, and mobile health apps: lessons from DSHEA and the regulation of dietary supplements,” *Maryland Law Review,* vol. 74, 2014.

[45] M. Boulos, A. Brewer, C. Karimkhani, D. Buller and R. Dellavalle, “Mobile medical and health apps: state of the art, concerns, regulatory control and certification.,” *Online Journal of Public Health Informatics.*

[46] J. Palmier-Claus and others, “The feasibility and validity of ambulatory self-report of psychotic symptoms using a smartphone software application,” *BMC Psychiatry,* vol. 12, no. 1, 2012.

[47] J. G. A. K. L.-P. M. GM Lucas, “It’s only a computer: virtual humans increase willingness to disclose,” *Comput Human Behav,* vol. 37, 2014.

[48] H. Joudaki and others, “Using data mining to detect health care fraud and abuse: a review of literature,” *Global journal of health science,* vol. 7, no. 1, 2015.

[49] S.-L. Wang and others, “The evaluation of trustworthiness to identify health insurance fraud in dentistry,” *Artificial intelligence in medicine,* vol. 75, 2017.

[50] I. Kose, M. Gokturk and K. Kilic, “An interactive machine-learning-based electronic fraud and abuse detection system in healthcare insurance,” *Applied Soft Computing,* vol. 36, 2015.

[51] J. Liu and others, "Graph analysis for detecting fraud, waste, and abuse in healthcare data," *AI Magazine,* vol. 37, no. 2, 2016.

[52] F. Araújo and others, "Using machine learning to support healthcare professionals in making preauthorisation decisions," *International journal of medical informatics,* vol. 94, 2016.

[53] R. Bauder and T. Khoshgoftaar, "A probabilistic programming approach for outlier detection in healthcare claims," in *Bauder, Richard A., and Taghi M. Khoshgoftaar. "A probabilistic programm15th IEEE International Conference on. IEEE*, 2015.

[54] D. Thornton and others, "Outlier-based Health Insurance Fraud Detection for US Medicaid Data," in *ICEIS (2)*, 2014.

[55] M. Johnson and N. Nagarur, "Multi-stage methodology to detect health insurance claim fraud," *Health Care Management Science ,* vol. 19, no. 3, 2016.

[56] R. Bauder, T. Khoshgoftaar and N. Seliya, "A survey on the state of healthcare upcoding fraud analysis and detection," *Bauder, Richard, Taghi M. Khoshgoftaar, and Naeem Seliya. "A survey on the Health Services and Outcomes Research Methodology,* vol. 17, no. 1, 2017.

[57] R. Bauder and T. Khoshgoftaar, "Medicare Fraud Detection Using Machine Learning Methods," in *Machine Learning and Applications (ICMLA), 16th IEEE International Conference on. IEEE*, 2017.

[58] D. Nicol and Otlowski, "Legal and Regulatory Issues of Precision Medicine.," Input paper for the Horizon Scanning Project "The Future of Precision Medicine in Australia" on behalf of the Australian Council of Learned Academies, www.acola.org.au., 2017.

[59] A. Turner, "The Future of Precision Medicine in Australia: Data.," Input paper for the Horizon Scanning Project "The Future of Precision Medicine in Australia" on behalf of the Australian Council of Learned Academies, www.acola.org.au., 2017.

[60] Productivity Commission, "Data Availability and Use," Canberra, 2017.

[61] OAIC (Office of the Australian Information Commissioner), "Community Attitudes to Privacy Report," Australian Government, 2013.

[62] Research Australia, "Australia Speaks! Research Australia Opinion Polling," 2016.

[63] The National Health and Medical Research Council, "National Statement on Ethical Conduct in Human Research," Canberra, 2007 (updated 2018).

[64] K. Tannous, "Case Study 1: The Real Cost of Accessing Linked Health Data Study".

[65] X. You and A. Okunade, "Income and Technology as Drivers of Australian Healthcare Expenditures," *Health Economics ,* vol. 26, no. 7, 2017.

[66] Australian Institute of Health and Welfare, "Health expenditure Australia 2015–16. HWE 68," AIHW, Canberra, 2017.

[67] Frost & Sullivan, "Transforming healthcare through artificial intelligence systems," 2016.

[68] Infosys, "Amplifying Human Potential: towards purposeful artificial intelligence," 2017.

[69] A. v. d. Heijden and others, "Validation of automated screening for referable diabetic retinopathy with the IDx-DR device in the Hoorn Diabetes Care System," *Acta ophthalmologica,* vol. 96, no. 1, 2018.

[70] T. Benson and G. Grieve, Principles of Health Interoperability, Springer, 2016.

[71] J. Mandel and others, "SMART on FHIR: a standards-based, interoperable apps platform for electronic health records," *Journal of the American Medical Informatics Association,* vol. 23, no. 5, 2016.

[72] A. Darcy, M. Alison, A. Louie and L. W. Roberts, "Machine learning and the profession of medicine," *JAMA ,* vol. 315, no. 6, 2016.

[73] J. Powles and H. Hodson, "Google DeepMind and healthcare in an age of algorithms," *Health and Technology,* vol. 7, no. 4, 2017.

[74] J. Chen and others, "The use of smartphone health apps and other mobile health (mHealth) technologies in dietetic practice: a three country study," *Journal of Human Nutrition and Dietetics,* vol. 30, no. 4, 2017.

[75] T. Patel and l. others, "Correlating mammographic and pathologic findings in clinical decision support using natural language processing and data mining methods," *Cancer,* vol. 123, no. 1, 2017.

[76] A. Hosny and others, "Artificial intelligence in radiology," *Nature Reviews Cancer,* vol. 1, 2018.