REPORT

A Before B

Perspectives on Artificial Intelligence in Healthcare

October 2018

CONTRIBUTORS:

- Angela Copeland, Cancer Care Ontario
- Zaki Hakim, University Health Network
- Darren Larsen, Ontario MD
- Kevin Jones, Strata Health
- Philip Barker, Fraser Health
- Peter Bascom, Formerly with eHealth Ontario

Digital Health
CANADA

Table of Contents

Executive Summary

The Canadian Health Informatics Executive Forum set out to help clarify the growing advancement in knowledge around Artificial Intelligence in healthcare. In order to accomplish this task a working group was initiated to:

- Establish a common understanding of Artificial Intelligence (AI)
- Identify the areas of AI relevant to the healthcare industry
- Conduct a literature and industry review of AI within healthcare
- Create some recommendations on how to prepare for AI within healthcare

Artificial Intelligence (AI) is typically defined as the ability of a machine to perform cognitive functions we associate with human minds such as perceiving, reasoning, learning, and problem solving. Examples of technologies that enable AI are robotics, and autonomous vehicles, computer vision, language, virtual agents and machine learning. This intelligence is brought about through some measure of learning. Within the healthcare space specifically there are three areas that this focusses on:

Natural Language Processing (NLP): this is the focus on having machines able to understand natural human language. Within healthcare this is important given the huge data sets created by clinicians in natural language. It is also integral to being able to interact with a human who is speaking vs directly interacting with forms.

Machine Learning (ML): this is the focus on having machines being able to learn from large data sets. Machine learning algorithms detect patterns and learn how to make predictions and recommendations by processing data and experiences rather than by receiving explicit programming instruction. This is relevant to healthcare in that all tasks that require the machine to be smart need to be able to learn from the data sets that exist within healthcare.

Deep Learning (DL): this is a subset of ML concerned with algorithms inspired by the structure and function of the brain called artificial neural networks. This approach uses neural networks to streamline the learning process vs having specifically coded rules. This is relevant to healthcare in the same manner as ML above.

There are a significant number of companies within the global and Canadian markets focussing on the adoption of AI within healthcare and the government of Canada is investing to make Canada a global AI hub. The healthcare industry and companies creating AI solutions still need to overcome the following barriers:

- Lack of funding to utilize solutions within the healthcare setting
- Lack of industry coordination to provide clear direction
- Lack of knowledge or understanding of the value proposition for AI in healthcare
- Access to authoritative unbiased big data sets
- Regulatory implications for the use of AI techniques

Organization executives need to prepare for AI as these barriers are broken down and can do so by:

- Acknowledging that AI starts with the data and data must be collected and made available for AI to learn
- Governance on who makes decisions about the introduction, investment and control of AI in health care organizations
- Understand and acknowledge the impact on health human resources, their roles and responsibilities, and qualifications in order to adopt and embrace emerging AI technology
- Think about AI in not just a clinical sense, but also in the context of the back office functions
- Consider the privacy and ethical considerations of AI as each potential new solution is reviewed
- Determine your role in AI and prepare your organization for a culture of change

There isn't much debate that AI will be a factor within healthcare and many see this change happening sooner rather than later. The confusion and lack of clarity around what AI is and what it can impact is becoming more and more clear. Executives within public and private sectors need to include AI in their planning and ensure that they are ready for the change that has already started happening.

There isn't much debate that AI will be a factor within healthcare and many see this change happening sooner rather than later.

What is Artificial Intelligence (AI)?

There is still a lot of confusion of what is Artificial Intelligence (AI) as confirmed in a recent survey of the membership of Digital Health Canada whereby most respondents felt they knew little to nothing about AI. For the purpose of this document the following definition in Wikipedia will be used:

> *Artificial intelligence (AI, also machine intelligence, MI) is Intelligence displayed by machines, in contrast with the natural intelligence (NI) displayed by humans and other animals. In computer science AI research is defined as the study of "intelligent agents": any device that perceives its environment and takes actions that maximize its chance of success at some goal. Colloquially, the term "artificial intelligence" is applied when a machine mimics "cognitive" functions that humans associate with other human minds, such as "learning" and "problem solving".*

If we take the above definition and think of Artificial Intelligence as the broader concept of machines being able to carry out tasks in a way that we would consider "smart", we can then think of machine learning (ML) as an application of AI based around the idea that we should really just be able to give machines access to data and let them learn for themselves.

We will focus on three major disciplines of AI and each has a definition below.

Natural Language Processing (NLP)

NLP is the AI focus of having computers process and analyze large amounts of natural language (human natural language). Specifically, the challenges being addressed are speech recognition, language understanding, and language generation. As humans we can say the same thing in many different ways and in many different languages; NLP must be able to unify this language into data constructs that can be processed.

Machine Learning (ML)

Puneet Gupta in the Harvard Science review describes machine learning as being split into three primary categories: supervised learning, unsupervised learning, and reinforcement learning. In supervised learning, a ML model is given data that has been labeled with a certain outcome, and then learns the relationship between both (data and outcome) to make predictions regarding the outcome for future data. In unsupervised learning, a ML model is given data that has not been labeled with an outcome, so it is able to sort and separate the data into groups of its choice, unlike supervised learning, which has certain outcomes or groups that the data must fit into. In contrast reinforced learning is learning by trial-and–error, in which a computer is instructed to achieve a stated goal in a dynamic environment. The program learns by repeatedly taking action and measuring the feedback for those actions and iteratively changing the behaviour until the desired goal is achieved. Overall, machine learning models attempt to adopt principles based on how humans innately learn and involves building systems that can 'think' and adapt themselves[32].

Deep Learning (DL)

Deep learning is subdiscipline of machine learning that focuses on learning data representations vs learning task-specific algorithms. The real power of deep learning is in how it can automatically learn representations without having to have coded rules. This is done through the use of multi-layered artificial neural networks.

Potential Benefits of AI

AI systems can provide advantages in speed, capacity, quality and consistency. It can handle repetitive tasks, processing large and complex data, and continuously monitoring operations. This suggests a significant opportunity to apply AI in healthcare based on the following:

- There are a significant number of repetitive tasks in the administration of healthcare
- There is significant variability in decisions and outcomes that lead to different treatment plans and outcomes, a step towards personalized medicine
- There is a vast amount of data in different formats that can be captured, analyzed, or leveraged
- There is untapped opportunity to make predictions of future outcomes based on current treatments and practices

This was confirmed in the Digital Health Canada survey, where a majority of respondents felt that it would definitely have an impact. The top four most appealing applications were as follows:

Patient/Client/ Person: Improved health outcomes, faster and/or more accurate diagnosis, and ideally more involvement in health care decision making.

Clinicians: Ability to personalize care paths and treatment plans, monitor medical conditions and/or health status and flag concerns, manage large amounts of data and diagnose conditions.

Administration: Predict capacity and system needs, mine medical records, optimize clinical processed to be better and faster.

Vendors: Ability to create intelligent workflows, ability to have technology self-learn, and ability to build technology that meets business requirements better.

Application of AI in Healthcare

Recent advances in analytics techniques, computational processing capabilities, and the increased availability of healthcare data has powered a paradigm shift in the status of Artificial Intelligence (AI) applications in healthcare. Some scientist even claim that AI may be able to replace all physicians, in the near future, by surpassing their logical and clinical decision making capabilities[1]. The availability of big data sets enables algorithms to refine their clinical decision-making capabilities and major advances have been made in utilizing these algorithms to improve healthcare outcomes by application in diagnostic imaging, bioinformatics/genomics, electrodiagnosis, and global screening[2].

Machine learning techniques have enabled scientists to cluster traits and enable identification of disease outcomes with increasing probability. These techniques include decision trees, naïve Bayes clustering, k-nearest neighbour, hidden Markov chains, logistic regression, neural networks, etc[2]. Natural language processing techniques have also improved our ability to extract information from unstructured data including, physician notes and journal articles, and integrated this information with structured data to improve clinical judgement[3].

Efforts are underway across the globe to utilize or invest in the development of AI techniques that mine for patterns in large electronic health care databases, and input those patterns into a mathematical framework. That framework then determines the best possible treatment for an individual patient. While significant literature exists on the academic achievements and breakthroughs in AI assisted healthcare, we will focus on real-world applications of these techniques.

Efforts are underway across the globe to utilize or invest in the development of AI techniques that mine for patterns in large electronic health care databases, and input those patterns into a mathematical framework.

I. Natural Language Processing (NLP)

NLP is essential in translating relevant medical knowledge, locked in unstructured medical notes, into structured data which can be utilized to build better models and computer processes to drive improvements in outcomes[4]. Today, a major portion of clinical observations for patients, including radiology reports, operative notes, and discharge summaries are recorded as narrative text either through dictation, transcription, or via direct entry into an EHR/EMR system[5].

NLP systems are able to represent clinical knowledge in standardized formats[5]. NLP is required in the clinical domain to power critical applications, such as decision support, cohort identification, patient management, resource management, question answering, knowledge acquisition, research, and discovery[4]. Within the broader biomedical domain NLP is required for applications such as information retrieval, database curation, knowledge discovery, knowledge acquisition and management, and tailoring information for consumers[4]. NLP systems are able to assess the quality of clinical information, monitor medication and diagnostic compliance, parse out and assign codes for disease identification, create structured data from unstructured data, generate and deliver patient-specific knowledge, and most importantly summarize and translate this information for patient use[5]. Ultimately the goal of NLP is to make data more accessible and actionable. Advances in linguistics and semantic capabilities of computer-based systems have allowed for the development of some exciting technologies.

Nuance for example, recently launched the Dragon Medical Virtual Assistant which features a cloud-based voice recognition technology that allows clinicians or patients, using specialised medical terminology, to converse with it in a natural manner with high accuracy[6]. This solution not only reduces the administrative workload for clinicians but also advances the patient education and engagement by allowing patients to query a smart machine to better understand their current health status.

Doc.ai is a platform which uses NLP to converse with patients via an app and explain their lab results to them[7]. If the communication with the patient reaches an impasse, the platform can refer patients to a human doctor. Doc.ai uses an edge-learning network to develop insights based on personal data, health records, wearable device data, and/or social media accounts[7]. The AI processes the information and starts drawing inferences between the datasets. Patients can ask questions to the app such as, "How can I decrease my cholesterol in the next 3 weeks?", or "Why was my glucose level over 100 and a week later it is at 93?" and receive responses in natural language[7].

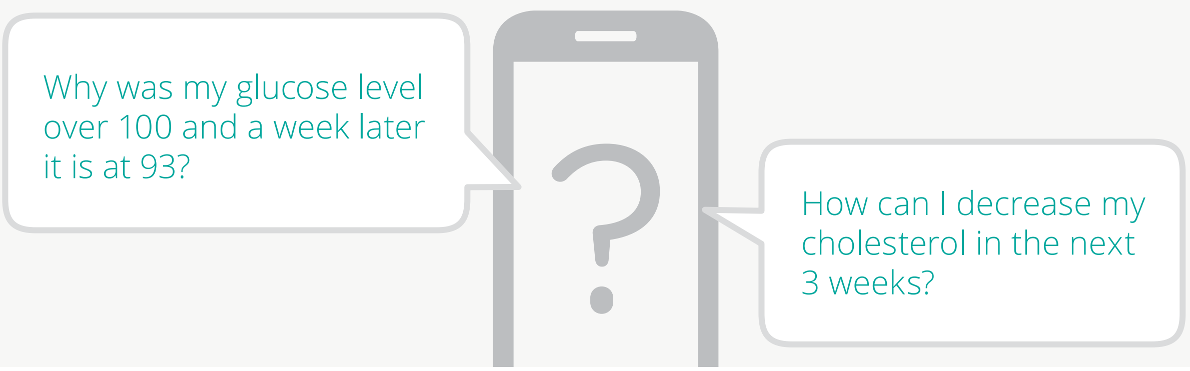

These examples showcase the ability of NLP to transform our ability to translate and communicate relevant health information with patients and improve the experience of clinicians by providing clinical decision support in a manner comparable to popular assistants such as Google Assistant, Siri, or Alexa.

II. Machine Learning (Supervised & Unsupervised)

Machine learning techniques are being leveraged around the globe to not only create systems that can perform at superhuman levels at Jeopardy but to also combat wicked problems in healthcare including population health management. Below we list some academic and real-world applications of AI in solving problems such as predicting readmission rates, improving screening rates for diseases, and combating cybersecurity threats.

Researchers were able use machine learning algorithms to improve the predictive ability of the instrument to classify autism with 100% accuracy

The Autism Diagnostic Observation Schedule-Generic (ADOS-G) is a commonly used instrument for evaluating the behavioral characteristics of autism spectrum disorders based on their language and developmental level[8]. Researchers were able use machine learning algorithms to improve the predictive ability of the instrument to classify autism with 100% accuracy[8]. The researchers were also able to accurately classify autism while winnowing 21 items from the ADOS-G thus improving the efficiency and speed of the test[8]. Results from the study will enable the development of mobile tools for preliminary evaluation and significantly speed the pace of initial evaluation and broaden our ability to identify the population at risk[8].

Researchers have also utilized supervised machine learning methods to mine medical/administrative information including patient's demographics, dispensed drugs, medical or surgical procedures performed, and medical diagnosis, to build predictive models for estimating the rate of readmission[9]. 21 classification experiments were performed (ROC, AUC, naïve Bayes classification, etc.) to model patient risk for readmission, highlighting the power of data science in utilizing a multitude of tools & methods to improve the accuracy of predictions[9].

An application of machine learning in surgical services includes Intuitive Surgical's da Vinci system which is helping surgeons by automating repetitive surgical subtasks such as suturing, cutting, and debridement[10]. Automating these surgical sub-tasks has presented challenges in the past due to the difficulty of modeling the precision required for cutting and operating on the highly nonlinear, viscous, and elastic nature of the human body[10]. However, researchers have been able to teach a surgical robot a finite set of surgical sub-tasks through a learning by observation technique wherein the robot practices a the surgical technique repeatedly and the underlying machine learning system refines the parameters of vision, technique, and execution time until the desired level of proficiency is achieved10(an example of reinforced learning).

Other examples of deployment of machine learning system includes Medicrea's UNiD system which uses machine learning to manufacture patient specific surgical implants to reduce the rate of post-operative failure and improve surgical outcomes[11]. Post-surgical spinal fractures were found to be reduced by 85% when the technology was utilized during surgery[11].

The benefits of machine learning are not limited only to clinical practice or the improvement of patient experience. Benefits may be realized in the form of improved cybersecurity. Cybersecurity systems that rely on supervised machine learning may fail in cases where the attacker 'fools' the system by teaching it a malicious pattern. Advances in unsupervised learning for cybersecurity now allows for the development of machines that do not depend on rules or signatures as it learns what is normal within a network[12]. The machine constantly revisits assumptions about behavior, using probabilistic mathematics without any reliance on human input[12]. DarkTrace, a cyber-defence company, utilizes recursive Bayesian estimation to constantly adapt the security settings of the system to enable it to better combat emerging cybersecurity threats without the need for training or instructions[12].

III. Deep Learning

Deep learning methods such as artificial neural networks (ANNs) offer a wide variety of advantages over supervised & unsupervised learning methods[13]. These include:

- Optimization ease, resulting in flexible and non-linear modelling of large data sets
- Improved accuracy of prediction inference
- Artefact detection with the ability to automatically re-estimate the model to provide improved clinical decision support

Artificial neural networks have been used in multiple studies to build models for clinical decision support in cancer, often with promising results[13]. 21 trials have found an increase in benefit to healthcare provision when ANNs were used in the diagnosis, prognosis, and therapeutic guidance for cancer[13]. Convolutional Neural Networks (CNNs) have been shown to be on par with 21 board-certified dermatologists in classifying biopsy-proven clinical images of different types of skin cancer including keratinocyte carcinomas and malignant melanomas[14]. Stronger than baseline classification of several diseases, including early prediction and diagnosis is possible, if a wide array of quality data sets is available to train the neural network models[14]. Some examples include early diagnosis of Alzheimer's from brain MRIs[15], early prediction of re-admission risk[16], diagnosis of breast nodules from ultrasound images[17], and prediction of congestive heart failure and chronic obstructive pulmonary disease from longitudinal EHR data[18].

Moving away from the healthy volume of academic successes in leveraging neural networks for medical uses, we also find an abundance of real-world applications. Enlitic, a deep learning artificial intelligence startup is using a lung cancer screening solution powered by a 3D deep learning engine[19]. This engine allows Enlitic to detect and characterize lungs nodules, longitudinally monitor diagnostic findings, and generate automatic reports to enhance sensitivity and scale in lung screening programs[19]. This reduces the need for needless biopsies, which are known to lead to significant volume of false positive diagnosis[19]. The solution has also been rolled out for thoracic imaging and digital mammography.

The cloud-based medical imaging company, Arterysis, has developed a cardiology application that uses AI to quantify blood flowing through the heart[20]. Their deep learning algorithms utilizes approximately 10 million rules and MRI images to produce contours of each of the heart's four chambers[20]. This allows the algorithm to precisely measure the volume of blood the chambers move with each contraction. Today, this calculation is manually completed by doctors and requires 30 to 60 minutes to measure the volume of blood transported with each pump[20]. In contrast, the algorithm comes up with the answer in 15 seconds! The solution has been cleared by the FDA and Arterysis has developed a partnership with GE Healthcare to incorporate the technology in the GE MRI machines. Improved scaling, partnerships, and regulatory approval will enable many similar applications to soon enable vastly superior diagnostic capabilities, powered by robust mathematical models, thus merging healthcare with the realms of mathematics, computer science, biotechnology, and statistics.

This allows the algorithm to precisely measure the volume of blood the chambers move with each contraction. Today, this calculation is manually completed by doctors and requires 30 to 60 minutes to measure the volume of blood transported with each pump. In contrast, the algorithm comes up with the answer in 15 seconds!

In April 2018 the FDA approved the first software powered by AI that replaces the need for a specialized doctor to interpret medical images. IDX-DR[21] is an automated detection software for diabetic retinopathy in primary care. As an autonomous, AI-based system, IDx-DR is unique in that it makes an assessment without the need for a clinician to also interpret the image or results, making it usable by health care providers who may not normally be involved in eye care.

AI Healthcare Advancements in Canada

Prominent AI researchers such as Geoffrey Hinton, Ph.D., Yoshua Bengio, Ph.D., hail from Canada and have laid the foundation for an AI boom in the country. Hinton is considered a pioneer in training neural networks with multiple layers while Bengio is a noted expert in deep learning[22]. The presence of heavyweight researchers and the recent investments from the federal government has catapulted Canada to be a hub for AI research. The foundation of the Vector Institute in Toronto is a significant opportunity to make Canada a global leader in artificial intelligence[22]. Google Deepmind and IBM have also established research efforts in Edmonton and Hamilton, respectively[23, 24]. These sandboxes will allow Canadian AI innovation efforts to flourish in the upcoming future.

As investment pours into the Canadian AI space, we will see an increased proliferation and utilization of AI technologies across various industries. Healthcare is also seeing its fair share of AI innovation. Winterlight Labs has created a machine learning solution which analyzes natural speech to detect and monitor dementia, Alzheimer's, aphasia[25]. The system uses short recorded speech samples to analyze linguistic cues, and can detect cognitive diseases with high accuracy (88-100%)[25]. Future use cases include training the system to recognize subtle change in speech to diagnose early onset of these cognitive diseases[25]. This represents a major improvement over current pencil-and-paper tests which are time-consuming and difficult to administer.

AI is also helping combat colorectal cancer, the second deadliest cancer in Canada. Artificial Intelligence in Medicine (AIM), a Toronto based company, has developed Rapid Case Ascertainment (RCA) software to read colonoscopy pathology reports using natural language processing[26]. The machine learning solutions reviews narrative text within the report and determines if the relevant medical concepts are expressed, allowing for the categorization of the number of patients undergoing the procedure, the adenoma detection rate (ADR), and polypectomy rate (PR)[26]. The system employs NLP to detect positive or negative context within the narrative to accurately determine the above features rather than counting the number of mentions[26]. This allows the system avoid counting the phrase "No signs of adenoma" when generating an ADR report. The AI solution thus dramatically increases efficiency of report generation, reduces human error, reduces cost associated with manually reviewing reports, and enables easily calculation of quality indicators for performance measurement at the individual, local, and provincial level[26].

A Nova Scotia based AI company is also joining the ranks with advanced breast imaging analytics. Densitas is utilising its breast density software to help radiologists provide personalized follow-up care to women[27]. Dense breasts can mask breast cancer as they are composed of glandular tissue which may resemble healthy tissue on a mammogram27. Breast density is also linked to an increased risk of having or developing cancer27. The AI solution determines breast density and informs clinicians to ensure that a patient with dense breasts returns for more frequent breast screening and follow-up care.

These examples represent a fraction of the innovative solutions that are being research, developed, and disseminated across Canada. According to Digital Health Canada survey respondents, Amazon Web Services and SAS are active in the vendor community, and in the public sector there is Alberta Health Services, Canadian Institutes for Health Research, and the establishment of the Centre for Excellence in Healthcare AI in Niagara Falls, Ontario. As funding, interest, and research grows within the AI space, a variety of novel solutions will undoubtedly become available.

The AI solution thus dramatically increases efficiency of report generation, reduces human error, reduces cost associated with manually reviewing reports, and enables easily calculation of quality indicators for performance measurement at the individual, local, and provincial level.

Preparing for AI

Due to the enormous buzz around AI solutions, extremely high expectations have been placed on the impact of AI in transforming how medicine is practiced and healthcare is delivered. We would like to begin this recommendations section by cautioning against treating AI solutions as a panacea. AI solutions, while driven by enormous quantities of data, do not always have sound logic built into the algorithms that may be utilized, which could lead to negative consequences when using these systems for clinical decision support. The old saying of "garbage in, garbage out" is almost weaponized by AI solutions because they rely completely on the data provided and develop their logic based exclusively on the quality of data provided, incorporating biases and incorrect information/assumptions.

With this caveat in mind, we would suggest that you keep the following recommendations in mind when developing your AI strategy:

1. To realize the power of artificial intelligence, start with data

It is very easy to get caught up in the hype of AI but before you start investing time and energy into AI with the intention to drive innovation be sure to assess your current environment and ensure you have a solid data foundation and governance program. If the foundation and practices are weak, the solution may not be optimized or worse completely off track. The first step is to think about your data collection practices. Are you collecting the right data to reach your ultimate goal of the type of AI investment you are thinking of making? Historically, healthcare has been a very paper-based industry therefore the big push was to collect any type of data and as much as possible. Advancements in technology has now allowed more selective practices focused on collecting data with the mindset of meaningful use that can be actioned resulting in better outcomes. However, are these the same data sets required for AI or do they need to be augmented with other data such as socio-economic, environmental, behavioural, genomic, etc.? And are those data sets from valid sources – can they be trusted? Another question to ask is does it need to be an on-going flow of real-time data and how much of it needs to be retained for the learning process? In other words, how much data is required to determine a repeatable pattern and how often does it need to be refreshed to detect deviations in the pattern that would influence the AI decision-making process to ensure accurate predictions. Finally, what level of quality is required to ensure that false predications do not occur and what level of cleansing and standardization is required to accelerate learning? These are all key questions that should be consistently asked by an organization considering investing in AI and should be embedded in data governance practices. It is critical to apply a new way of thinking about data governance that is more agile in that it can handle the speed of change, learn and adapt quickly from success and failure, be more democratic while ensuring the protection of individual and collective rights and interests. If you are confident that you can collect the needed real-time data, that it is organized, clean, tested and optimized, then it is time to test machine learning and artificial intelligence solutions.

2. Don't exclusively focus on the clinical applications, instead shift your attention to automating back-office functions with AI

Process that traditionally relied on labour-intensive manual interventions will see greater consistency, better controls, and improved quality once they are automated using AI-driven solution. This frees up knowledge talent within the organization to perform knowledge tasks that add greater value, such as questioning, refining, and learning. Since most back-office functions have already been automated within the finance, manufacturing, and retail industries, automating these processes represent quick-wins for AI-enabled solution within the healthcare ecosystem. The efficiencies realized from freeing up knowledge talent that has previously been occupied with manual interventions such as contract management, reporting, technical help, etc. may be the key to driving improved innovation efforts within the healthcare sector.

Examples of the value of AI solutions in the back-office includes, HyperScience, a firm that leverages computer vision and image recognition technologies to scan forms, recognize handwriting with higher-than-human accuracy, and automatically fill in all relevant data fields within an enterprise system[28]. The system can parse through human language in a claim, determine relevant or irrelevant information, and highlight all the information needed to make correct assessments and decisions[28]. Such a solution can deliver tremendous value for both health service providers and their customers (patients, policy holders, applicants) by expediting the resolution period for insurance claims, medical treatment, etc.

Other relevant examples include, automation of procurement activities. Machine learning solutions can quickly scan contract data, and extract important information from a contract including names, organization and vendor information, the contract signature date, renewal dates to ensure that contract end dates are never missed and cost savings are realized by exploiting relevant clauses[29].

Implementing back-office AI solutions will be made easier by the fact that Financial institutions, Telecom giants, and various other industries already possess the subject matter expertise and strategies to successfully navigate these solutions. Knowledge sharing through partnerships and collaborations with established industries will ensure that the initial forays into AI solutions within the Canadian healthcare space are a rousing success. Learnings from these successes may then be leveraged to design better solutions within the clinical decision support space. A focus on back-office AI solutions offer quick-wins, cross-industry partnerships, and learnings to build up the AI skillset to succeed within a more complex environment.

3. Pay close attention to the privacy, legal, and ethical challenges associated with AI solutions

AI solutions are capable of tracking and predicting shopping, political, and location preferences. As the availability of rich structured and unstructured data expands, it will become increasingly easier to pinpoint individual's based on preferences, medical history, or location[30]. This may limit our ability to appropriately de-identify patient data and negatively impact freely available data for clinical research. We must be careful to ensure that de-identification solutions are developed in tandem with the global surge in research to pinpoint network and individual preferences for advertising purposes.

On the ethical side, there is a persistent belief that artificial intelligence will be able to perform more objective analysis and thus produce objective solutions. However, algorithms and models are as objective as the people who build them and the data that is used for their training. The model's result may be incorrect and even discriminatory if the training data renders a biased picture of reality or if the relevance of data is in question. Medical datasets openly available for use by AI researchers are notoriously biased towards the white male populations. A 2014 study tracked cancer mortality over 20 years and found that a lack of diverse research subjects was a key reason in why black Americans are significantly more likely to die from cancer[31]. AI solutions would incorporate these inherent biases which are hidden within large data sets and may produce results that may prove harmful or discriminatory against some populations. As such, we must require that AI models be trained using relevant and correct data. Furthermore, models must be designed so that they do not emphasize information relating to ethnic origin, political opinion, religion or belief, genetic status, health status, or sexual orientation if this would lead to arbitrary discriminatory treatment[32].

Finally, from a legal standpoint, it is necessary to ensure that algorithms can make predictions that are traceable or interpretable. Black-box solutions, mostly prevalent in deep learning techniques, are not only opaque but change over time as the data guiding their training changes[33].

> …Algorithms can make predictions that are traceable or interpretable. Black-box solutions…are not only opaque but change over time as the data guiding their training changes.

4. Determine your role in the AI ecosystem and create a culture of AI

AI is here. The question will not be whether or not you participate it will be how do you participate, where do you invest, and what does that investment look like. AI will be most successful in healthcare if there is a trusted network of organizations that work together to “develop”, “curate”, and “apply” to realize the benefits. “Development” of AI solutions is a significant financial investment to ensure that there is enough computing power, data, expertise, and tolerance for failure to deliver results. It is ideal if this can predominantly happen in the research and vendor community. “Feed” is the accountability to produce good quality data whether you are a healthcare provider, clinician or even a patient. By investing in creating a culture of good data practices the better the downstream AI outcomes will be. “Apply” is the fostering the culture to embrace AI and use it in practice. Interestingly, according to Digital Health survey respondents, the top concern regarding AI was humans being too reliant on AI and the lowest concern was AI machines becoming more intelligent than humans. Overall a majority of respondents were excited and optimistic about the prospect of an AI future in healthcare. However, this is not the norm as there is still a lot of apprehension and skepticism around how AI will have an impact. It is key to debunk myths now. AI will never replace human reasoning and intuition – it is meant to augment the critical thinking process. Therefore, more effort should be placed on preparing people on how to use AI effectively starting with a formalized change readiness program. The program should start with establishing a common understanding of AI and conveying the benefits that come with using AI throughout the organization. It should also focus on providing people with the skills and know-how of how to effectively apply AI in day-to-day decision-making. By determining, what role your organization wants to have in AI it allows focus on building the required capabilities and establishing the required partners so overall we as a system can realize the benefits of AI.

Conclusion

84% of survey respondents felt it would take 0-5 years for AI to have a noticeable impact on the healthcare technology market

Though there is much confusion and debate about what AI actually is and what specific impacts it will have on healthcare everyone agrees: AI will have an impact on healthcare in the short term. This was confirmed within the Digital Health Canada survey where 84% of survey respondents felt it would take 0-5 years for AI to have a noticeable impact on the healthcare technology market and that overall technology would be better because of AI.

This is also clearly demonstrated by the growing number of companies releasing products and conducting research not just globally, but within Canada as well. These companies are actively working to address the barriers highlighted by the Digital Health Canada survey as lack of both funding and market coordination.

Furthermore, regulatory bodies like the FDA are recognizing that AI holds enormous promise for the future of medicine and are therefore, investigating developing new regulatory frameworks to promote and support the use of AI technologies.

The use of AI in the healthcare system is already here. Ensuring that your organization is ready to embrace this technology is crucial. AI is poised to make medical equipment smarter, patient examinations more precise and to be an enabler in personalized medicine. Therefore, as these products begin to impact the industry stakeholders need to make data sets available, identify areas that can benefit from AI, embrace and create a culture of change, while paying attention to any privacy and ethical considerations that this change creates.

References

1. Darcy, A. M., Louie, A. K., & Roberts, L. W. (2016). Machine learning and the profession of medicine. *Jama*, 315(6), 551-552.

2. Jiang, F., Jiang, Y., Zhi, H., Dong, Y., Li, H., Ma, S., ... & Wang, Y. (2017). Artificial intelligence in healthcare: past, present and future. *Stroke and Vascular Neurology*, svn-2017.

3. Murff, H. J., FitzHenry, F., Matheny, M. E., Gentry, N., Kotter, K. L., Crimin, K., ... & Speroff, T. (2011). Automated identification of postoperative complications within an electronic medical record using natural language processing. Jama, 306(8), 848-855.

4. Friedman, C., Rindflesch, T. C., & Corn, M. (2013). Natural language processing: state of the art and prospects for significant progress, a workshop sponsored by the National Library of Medicine. *Journal of biomedical informatics*, 46(5), 765-773.

5. Demner-Fushman, D., Chapman, W. W., & McDonald, C. J. (2009). What can natural language processing do for clinical decision support?. *Journal of biomedical informatics*, 42(5), 760-772.

6. Nuance (n.d.). Nuance Unveils Virtual Assistant Designed for Healthcare Providers. Retrieved December 14, 2017, from https://www.nuance.com/about-us/newsroom/press-releases/nuance-unveils-virtual-assistant-for-healthcare.html

7. De Brouwer, W., Mason,B. (2017). NeuRoN: Decentralized Artificial Intelligence, Distributing Deep Learning to the Edge of the Network. Doc.ai

8. Wall, D. P., Kosmicki, J., Deluca, T. F., Harstad, E., & Fusaro, V. A. (2012). Use of machine learning to shorten observation-based screening and diagnosis of autism. *Translational psychiatry*, 2(4), e100.

9. Hosseinzadeh, A., Izadi, M. T., Verma, A., Precup, D., & Buckeridge, D. L. (2013, July). Assessing the Predictability of Hospital Readmission Using Machine Learning. In IAAI.

10. Murali, A., Sen, S., Kehoe, B., Garg, A., McFarland, S., Patil, S., ... & Goldberg, K. (2015, May). Learning by observation for surgical subtasks: Multilateral cutting of 3d viscoelastic and 2d orthotropic tissue phantoms. In Robotics and Automation (ICRA), *2015 IEEE International Conference on* (pp. 1202-1209). IEEE.

11. A. Vaccaro, V. Fiere, S. Fuentes, T. Raabe, P. Passias, T. Protopsaltis, A. Faure, P. Tropiano, B. Blondel. (2017). Patient-Specific Rods show a reduction in rod breakage incidence. Medicrea.

12. DarkTrace. (2016). Machine Learning: A Higher Level of Automation. DarkTrace Limited.

13. Lisboa, P. J., & Taktak, A. F. (2006). The use of artificial neural networks in decision support in cancer: a systematic review. *Neural networks*, 19(4), 408-415.

14. Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639), 115-118.

15. Liu, S., Liu, S., Cai, W., Pujol, S., Kikinis, R., & Feng, D. (2014, April). Early diagnosis of Alzheimer's disease with deep learning. In *Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on* (pp. 1015-1018). IEEE.

16. Nguyen, P., Tran, T., Wickramasinghe, N., & Venkatesh, S. (2017). $\mathtt {Deepr} $: A Convolutional Net for Medical Records. *IEEE journal of biomedical and health informatics, 21*(1), 22-30.

17. Cheng, J. Z., Ni, D., Chou, Y. H., Qin, J., Tiu, C. M., Chang, Y. C., ... & Chen, C. M. (2016). Computer-aided diagnosis with deep learning architecture: applications to breast lesions in US images and pulmonary nodules in CT scans. *Scientific reports*, 6, 24454.

18. Cheng, Y., Wang, F., Zhang, P., & Hu, J. (2016, June). Risk prediction with electronic health records: A deep learning approach. In *Proceedings of the 2016 SIAM International Conference on Data Mining* (pp. 432-440). Society for Industrial and Applied Mathematics.

19. Enlitic. (n.d.). Enlitic Brings Deep Learning Diagnostic Solutions to the Radiological Society of North America. Retrieved December 14, 2017, from https://www.enlitic.com/press-release-11162016.html

20. Molteni, M. (2017, June 03). If You Look at X-Rays or Moles for a Living, AI Is Coming for Your Job. Retrieved December 14, 2017, from https://www.wired.com/2017/01/look-x-rays-moles-living-ai-coming-job/

21. IDx (2018) FDA permits marketing of IDx-DR for automated detection of diabetic retinopathy in primary care. Retrieved from https://www.eyediagnosis.net/single-post/2018/04/12/FDA-permits-marketing-of-IDx-DR-for-automated-detection-of-diabetic-retinopathy-in-primary-care

22. Research & Development. (2018). Why Canada is Becoming a Hub for A.I. Research. Retrieved from https://www.rdmag.com/article/2017/07/why-canada-becoming-hub-ai-research

23. Vanian, J. (2018). Google's DeepMind Turns to Canada for Artificial Intelligence Boost. Fortune. Retrieved from http://fortune.com/2017/07/05/google-deepmind-artificial-intelligence-canada/.

24. Ibm.com. (2018). IBM - Hamilton looks to AI to transform its healthcare system - Canada. Retrieved from https://www.ibm.com/ibm/ca/en/gm-hamilton-ai-helps-transform-health-care-system.html.

25. Fraser, K. C., Meltzer, J. A., & Rudzicz, F. (2016). Linguistic features identify Alzheimer's disease in narrative speech. Journal of Alzheimer's Disease, 49(2), 407-422.

26. Golabek, J. (2017). Using Natural Language Processing to Monitor Quality in Colorectal Cancer Screening – Artificial Intelligence in Medicine Inc. Retrieved from http://www.aim.ca/colorectal-cancer-screening/

27. Innovacorp. (2018). Densitas Rolls Out Machine Learning Innovation for Breast Health Across Canada. Retrieved from https://innovacorp.ca/news/densitas-rolls-out-machine-learning-innovation-breast-health-across-canada.

28. Yao, M., Zhou, A., & Jia, M. (2018). *Applied Artificial Intelligence: A Handbook For Business Leaders*. TopBots.

29. Lhoumeau, M. (2017, October 12). AI, Machine Learning and What They Really Mean for Procurement and Contracts. Retrieved from http://spendmatters.com/2017/10/12/ai-machine-learning-really-mean-procurement-contracts/

30. AI NOW, AI Now 2017 Report, 2017, available from: https://ainowinstitute.org/AI_Now_2017_ Report.pdf.

31. Aizer, A. A., Wilhite, T. J., Chen, M. H., Graham, P. L., Choueiri, T. K., Hoffman, K. E., ... & Nguyen, P. L. (2014). Lack of reduction in racial disparities in cancer specific mortality over a 20 year period. Cancer, 120(10), 1532-1539.

32. Datatilsynet. (2018). Artificial intelligence and privacy. Retrieved September 9, 2018, from https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf

33. Puneet Gufta, (2017). Machine Learning: The Future of Healthcare. Harvard Science Review. https://harvardsciencereview.com/2017/05/16/machine-learning-the-future-of-healthcare/

Digital Health Canada connects, inspires, and educates the digital health professionals creating the future of health in Canada. Our members are a diverse community of accomplished, influential professionals working to make a difference in advancing healthcare through information technology. Digital Health Canada fosters network growth and connection; brings together ideas from multiple segments for incubation and advocacy; supports members through professional development at the individual and organizational level; and advocates for the Canadian digital health industry.

For more information, visit **digitalhealthcanada.com.**

Digital Health CANADA

1100–151 Yonge Street, Toronto, ON M5C 2W7

647.775.8555

info@digitalhealthcanada.com