The Leading Edge of Disruptive Innovations in Health Care Data, Communications, and Information Technologies:

Artificial Intelligence

An Early 2018 Primer and Perspective

Gregory Downing, D.O., Ph.D.

Produced for the Global Digital Health Partnership, 2018
NOT FOR DISTRIBUTION

ABSTRACT

This issue brief explores the topic of artificial intelligence (AI) and the basis for its current surge in broad interest across the health care sector. Here, we explore AI from a historical perspective, terms and definitions, the major science and technology components, and, relevant to the Global Digital Health Partnership (GDHP), the newfound integration of artificial intelligence with health care and government. Our intent is to provide a general view of the pressing issues that are emerging at a rapid pace with this fascinating technology, and our work is not intended to represent a thorough, in-depth analysis from a policy perspective on science and technology implications. In addition, examples of the application of AI in consumer health, medical practice, and health care management are provided to lay the groundwork for conversations about the wide-ranging potential for these technologies to have a powerful impact and societal benefit. As a result of the basic understanding presented here, coupled with unique perspectives and experiences of each of the individuals in the audience, we encourage thoughtful introspection and commentary to be followed by considerations for collaborative activities on the topic among the GDHP participants.

INTRODUCTION

The applicability of machine learning (ML) and artificial intelligence (AI) technologies to health care is relevant and imminent with the long-term prospects anticipated to be extensive in many aspects. The implications of introducing intelligent systems into health care settings present opportunity for incredible advancements for research, clinical services, public health, health systems design and management, consumer health and beyond. Some of those advances are already being realized (e.g., medical imaging augmentation), some are on the horizon (e.g., personalized medicine with applications in predicting individual response to therapies based on genetic or other biological characteristics of disease states), and some at this point are speculative (e.g., automatic diagnosis and precision treatment plans), albeit possible.

Here, perspective is placed on contemporary issues related to the uses of AI in civic society for the purposes of framing considerations that provoke implications and considerations more narrowly in health care including at a level of ethical, economic, legal, and sociological dimensions. We begin by examining the definition(s) of AI as a decomposition to current methodologies and technologies. This is important as the term “AI” has come to represent a variety of concepts from its scientific applications in analytics, to those of the lay public and integration into the activities of their daily lives. These include supervised and unsupervised machine learning, convolutional neural networks, and reinforced learning.

We explore current trends and examine the potential opportunities and applications of AI and ML in health and health care. This is not intended to answer technical or methods questions, resolve issues of hyperbole, or prognosticate about the economic or workforce impact. However, suitable, this is a starting point for considerations for how arguments are presented with accompanying challenges and caution as to what factors have or could inhibit adoption or serve as severe obstacles to the technology's success. Concepts are presented with the intent of supporting discussion that may evoke the considerations by governments and civic organizations in the study of implications of this science and technology beyond solving engineering, and mathematical problems.

The ability that consumers, patients, and caregivers now have to hold smart technologies in their hands has empowered them to be more informed, engaged, and decisive about their care and purchasing. Given these advances, AI enables autonomy for the individual in the form of products and services that support decision-making in health and wellness concerns with independence from medical care providers. That same technology is able to interpret more about how humans live their lives which consequently is then represented as data for others to use. In health care and medicine, these data are being derived from new digital sources like electronic health records, medical devices, and sensors that enable the linkage to personally generated health data that can provide medical insights. In the future, it is likely that these personally derived data streams will converge with new enabling technologies such as virtual reality, autonomous mobility devices, exoskeleton, and neurocognitive enhancement technologies.

In many ways, these technology drivers have led to the current "big data" era and now serve as the antecedents to powerful analytic approaches that represent an element of the fuel needed to create "smart systems" of health care. These data include everything from geospatial coordinates to medical images, to entire genome sequencing data, that when coupled together can represent new medical intervention strategies that could have been imagined only a short time ago. The coupling of the massive growth in data generation and advanced analytic skills are catalyzing the growth of predictive analytics and modeling in population health, health systems research, and health economics. Taken together, these vast bodies of data and the analytic capability that AI is presenting have the potential to dramatically improve the value and productivity of health care services.

ARTIFICIAL INTELLIGENCE DEFINED

For the purposes of the topics discussed herein, artificial intelligence or "AI" is defined as, "the science and engineering of making intelligent machines" as it was in 1956 by John McCarthy, who coined the phrase. While often spoken of in the same context as AI, 'Machine Learning' is an application toward, or an approach to, artificial intelligence. Machine Learning (ML) is a method of using algorithms to train machines, or have them train themselves, to a point that the machine can make accurate decisions and predictions on its own. Recent history has allowed machines to solve problems that humans present as input to be computed, when the machine is programed to specifically solve those problems. Machines have also been able to alter the way information is portrayed to allow humans to detect trends and make predictions.

Two recent developments in the technological space have brought on the increased applicability and visibility of AI: 1) the computational power of machines has grown exponentially in the past two decades, so much that machines can handle much greater calculations than ever before and, 2) the development of machine learning techniques (e.g., deep neural networks) combined with the computational power mentioned, has given machines the ability to identify and solve several problems from one set of inputs, determine trends, make predictions, and even recommend decisions. While recent demonstrations have shown these actions can sometimes be made with levels of accuracy equal to or greater than the accuracy levels at which humans make the same determinations, many false positives still arise in machine learning outputs. If gone undetected, these variances introduce a bias in machine learning based on the training sets. The need for high quality data is vital to training machines to make the right decisions.

A BRIEF HISTORY OF ARTIFICIAL INTELLIGENCE

The history of AI is one for argument and arbitrary attribution, mostly because the science is composed of multiple theories and applications. Briefly, the term is drawn from the applications in the field of psychology and is often cited to 1943 with McCulloch and Pitts' publishing of "*A Logical Calculus of the Ideas Immanent in Nervous Activity*," which created early frameworks of today's neural networks.

It has also been discussed that AI's history began with Alan Turing during his quest to create intelligent machines, and then the publishing of his Turing Test (1950). The Turing Test was proposed as a sequence of questions asked of a party to determine if that party is human or not (i.e., a machine). These historic milestones are important to understand in today's context, as well as in the context in which they were introduced. The rate of AI's growth is similar to, and reliant on, the growth of computational power and predictive algorithms. Major developments in AI have occurred more often in recent history due to major increases in the computational power of Graphics Processing Units and the rapid development and accessibility of algorithms to drive more accurate or 'smarter' AI. The development and integration of physical intelligent systems did not begin until the late 1970's, while theory and publications thereof had existed for almost three decades prior. However, once systems began to be developed, the rate at which they were introduced grew as well. Beginning with simpler game-based approaches, like MIT's MacHack winning chess tournaments, development progressed to speech recognition and translation by IBM in 1988, IBM's Deep Blue chess champion in 1997, and ultimately Google's AlphaGo defeating the world's best Go player, Lee Sedol, in 2016. A more detailed overview of the historical timeline of advances in AI is included in the appendix.

As groundbreaking as each milestone may be, there have been setbacks that have stalled adoption and created doubt in the success of AI implementation. Setbacks can be as subtle as movies portraying negative influences of artificially intelligent systems, as mediocre as celebrities voicing their concerns of possible dire outcomes, and as actual and major as attempts that truly fail to meet their goals. This last example is best represented in recent health care history with the IBM Watson and MD Anderson Cancer Center 2013 partnership, using a, "cognitive computing system for its mission to eradicate cancer," (Herper, 2017). That project faltered in late 2016, which followed complications in IBM Watson's partnership with Memorial Sloan Kettering. "An investigation has found that the supercomputer isn't living up to the lofty expectations IBM created for it. It is still struggling with the basic step of learning about different forms of cancer. Only a few dozen hospitals have adopted the system, which is a long way from IBM's goal," (Ross and Swetlitz, 2017).

While impediments have delayed adoption and skewed some views of the technology, AI remains a growing industry and being democratized to the public. Every few weeks, new online courses emerge to teach any willing user how to immerse in deep learning, algorithm application, and data science. The ease of access to low cost cloud computing services has fueled the rapid growth of independent research and development by not constraining its development applications to large computing centers and corporate structures. In mid-2017, CB Insights reported that investments in, "AI in health care startups increased 29% year-over-year to hit 88 deals in 2016, and are already on track to reach a 6-year-high in 2017." The historic setbacks will likely become cautionary tales for new developers and organizations branching into AI. These concerns have prompted the emergence of new areas of scientific discourse including data ethics, workforce implications, and behavioral and psychological studies of its impact on human to human interactions. Speculation exists as to whether we are currently in another phase of AI popularity or if the technology will yield disruptive innovation throughout health care.

MACHINE LEARNING

A fundamental component of the current approach to artificial intelligence development is machine learning. With machine learning, machines are "taught" to learn the way humans learn, with examples and repetition. Any skill we have, be it cooking or identifying trends in data, the more we perform that task, the more consistently accurate our results become. The major difference is a machine's computational power and focus. A machine can intake thousands and millions of examples without once becoming tired or distracted. For example, a UK-based service, Your.MD, has been approved by the National Health Service (NHS) to use machine learning and artificial intelligence to augment the provision of health care information to patients (Your.MD, 2018). The idea is to provide "pre-primary care" and reduce the need for medical visits. The system uses user-input information against machine-trained algorithms. In one form of machine learning, Supervised Learning, examples are given to the machine depicting positive and negative results, and are labeled as such. The algorithm the machine is using to identify the patterns is what enables the machine's reasoning ability to then make predictions on images that are not labeled. This process is also known as classification. A common current example in health care exists in feature recognition such as that of medical imaging. With only 1000 images, of patient data, a machine can be trained to correctly classify cells from a pathology specimen of a tumor that have cancerous anatomic features with cancerous cells to a 97% accuracy when compared with those of expert human interpretation (CAMELYON17, 2017). An important distinction with supervised learning is that often machines are not much more accurate than the training examples, signifying the need for the best possible samples for training algorithms.

The next most utilized method of machine learning is unsupervised, which consists of clustering or associations. In this method, the preliminary, labeled examples do not exist. A machine is not given a classified test set from which to learn. It is given only a set of data and an algorithm and instructed to determine patterns and to group similarities it finds, producing an output to be analyzed and possibly categorized by humans. Clustering involves grouping patterns of like data, and association involves pairing data that present themselves in repetition. A current example of unsupervised learning exists today in image recognition. Not to be confused with the previously mentioned medical imaging example, machine-learned image recognition means determining unlabeled patterns in images, such as the repetition of a particular person's face in a set of images, or the ability of a machine to show images of playgrounds without a human instructing (e.g., coding) that machine to know what a playground is and is made up of.

NEURAL NETWORKS

A deeper application of unsupervised machine learning is artificial neural networks. Google's deep learning arm based in London, DeepMind, has demonstrated several successes in implementing artificial neural networks, including a recent application that reduced energy consumption in Google Data Center cooling by 40% greater than their engineers were able to design (Gao, 2016). In the whitepaper produced by Google explaining the application, neural networks are described as, "A class of machine learning algorithms that mimic cognitive behavior via interactions between artificial neurons." The concept is such that machines can achieve a greater depth of knowledge by utilizing an artificial network of neurons, where each artificial neuron is a mathematical function (e.g., sigmoid) and the collective solution of those neurons provides a more accurate and robust response than applications designed to produce a single output from a single input (Goodfellow et al, 2016).

Another noteworthy comparison to explore is deep learning neural networks versus reinforced learning. Artificial neural networks create groupings and associations based on data. Reinforced learning is similar, except that at the completion of the machine running its procedure, the machine is given a simple positive or negative input based on its actions. The best way to think of reinforced learning is through games in which a series of actions with a large number of permutations are undertaken to solve a problem or reach a desired outcome. For demonstrative purposes in the development of AI, this has featured the game 'Go'. Go is an ancient Chinese board game with more possible configurations than atoms in the known universe ($> 10^{170}$). In 2016, DeepMind's AlphaGo defeated the world's best Go player, Lee Sedol, four games to one. As profound as this achievement was, one year later, AlphaGo was defeated 100 games to zero, by DeepMind's newer machine, AlphaGo Zero (Silver et al, 2017). The premise undertaken in the original AlphaGo experiment, and even more so with AlphaGo Zero, was the idea of reinforced learning, where the machine attempts a combination of moves or procedures and is given feedback as to whether those attempts produced positive or negative results. The machine then refines its strategy repeatedly in many multiple series of attempts and improves the outcomes and chances of winning with each successive iteration.

ARTIFICIAL INTELLIGENCE IN HEALTH CARE

Artificially intelligent systems are now becoming more integrated into health care. Their initial applicability has been realized, but so have the challenges of adopting them. In addition to the increase in computational power previously discussed, a massive amount of digital health data are becoming available to not only health systems and providers, but to consumers as well. Health-related mobile applications now allow for the tracking of physical activity, both in duration and intensity. So-called “diary” apps allow for users to input their eating habits that have numerical attributes (e.g., calories). The location of potential medical patients is continuously tracked at a much higher rate than census can maintain. All of these exist as data outside of official medical records and create opportunities for new approaches to health and wellness in the marketplace. It seems to follow logic that intelligent systems that can help decide when cancer may be occurring based on the location of cells within an image could also help consumers decide on better food choices based on their location in the supermarket.

Machine learning can be utilized in feature recognition, particularly in visual and auditory representations of data. It is noted though that the application of such technologies is dramatically on the rise. A review of an organization dedicated to hosting international image analysis challenge competitions shows an increase from 15 to 24 challenges, from 2012 to 2017 (Consortium for Open Medical Image Computing, 2018); already in January of 2018, five challenge competitions are posted. Medical imaging is an area where major advances continue to be made. It is predicted here that the use of AI in medical imaging will drastically augment future radiological operations. Augmentation however, does not imply replacement. It is also foreseen here that augmentation of medical imaging will lead to better, safer imaging. Rote tasks will be supplanted, to be replaced with radiological experts trained to higher levels, able to detect and diagnose rare cases. Other areas of health care that are already seeing the application of artificial intelligence include natural language processing in electronic health records. The U.S. Department of Veterans Affairs conducted a study that classified correctly 100% of cases of heart failure using machine learning and natural language processing (Garvin et al, 2018). The areas in which AI is already seeing an impact will continue to grow, and newer technologies branching into further reaches of health will be explored as well.

One study completed recently is of noteworthy significance in its detail and utility. The MITRE Corporation’s advisory group, JASON, produced a report in late 2017 that evaluates artificial intelligence as an application to health and health care with regard to diagnostics, data collection, algorithm development, and challenges that may impede success, including legacy health records. With regard to medical imaging, the diagnosis of disease and conditions through imaging is carried out by specialists trained to detect anomalies versus normalcy (i.e., sensitivity vs. specificity) in situations where accuracy is vital to proper treatment. In the JASON study, retinal images were used to detect early signs of diabetes and deep learning algorithms were able to detect adverse anomalies with an accuracy (i.e. specificity) of 90-98%, which was as or more accurate than the training examples given in the supervised machine learning preparation. With regard to electronic health records, the study also reports on the United Kingdom’s ineffective diagnosis of cardiovascular disease due to redundant indicators combined with a lack of transparency in the AI algorithms:

> When the relationship between the data and the diagnosis suffers from error, variability or difficulty in discrimination, AI algorithms also perform less well. This creates challenges for developing AI for assessments based on data from EHRs, and foreshadows the opportunities (and challenges) of supplementing EHRs with extended patient reported data (JASON, 2017).

The JASON report concludes with notable recommendations, including assurances of privacy and transparency in applications and policies, sound data infrastructures to support AI, developing and supporting further research, and supporting competitions and challenges to further understand health and health care data.

HEALTH CARE AND AI ON THE HORIZON

Already on the horizon of health care systems is the introduction of truly personalized medicine – mass customization of health care interventions based on scientific evidence and personal preferences. Precisely defined medical strategies will be tailored specifically to patients based on intelligent analysis of their DNA or other biological markers, environmental influences, and medical history. The more data are introduced on reactions to drugs by genotype, the clearer indications will be to providers which drugs and treatments are more effective than others. The ideal vision allows not only for DNA and medical history to be factors in treatment, but a patient's social determinants, data from their mobile activity trackers, and a blueprint of their microbiome for reactionary indicators to drugs would all be orchestrated to optimize the treatment plan for maximum beneficial outcomes (Orion Health, 2017).

Another developing innovation in AI is drug discovery. Recent investigations have allowed researchers to map molecules in potential drugs to train machines on toxicity and side effect outcomes (Altae-Tran et al, 2017). In their test cases, the machines were able to better predict against adverse outcomes, which could lead to new more effective drugs. Breakthroughs in medicine using artificial intelligence is not possible without masses of high-quality data. In the field of drug discovery, projects like The Connectivity Map are collecting and standardizing these data into libraries with goals like uncovering, "novel treatments for a variety of diseases, including cancers, neurological diseases, and infectious diseases," (Clue.io, 2018).

While the future of health care will benefit greatly from customized technologies to revolutionize diagnosis and treatment, artificially intelligent technologies already exist today, that have not yet been adopted in health care, due to resistance, lack of a strategy for proper integration, and/or lack of a technology infrastructure that allows for the computational power needed to implement artificially intelligent systems. These technologies exist in speech recognition and automation. A 2014 health care study found a 97.3% accuracy in male speech recognition (Johnson et al, 2014). While consumerism has expanded for speech recognition devices, research reports that

50% of physicians are resistant to the technology despite the efficiencies and time savings the technology offers (KLAS, 2015). Lastly, automation of health care processes stands to provide drastic improvements in efficiencies and patient safety benefits, but adoption in health care services has been slow in the majority of health care delivery settings. Automation of patient intake processes, using intelligent systems to perform basic tasks such as logistics, scheduling, and alerting of personnel of abnormalities for prioritization are examples of steps that can fairly easily be taken, but have yet to be seen on a large scale in health care. Thus far, a fairly opportunistic picture has been painted of the health care industry with regard to the introduction of AI and ML solutions. However, many obstacles exist currently and will continue to arise inhibiting adoption and impeding real revolution in how health care is delivered.

Despite the promise and opportunity for AI, there are key questions emerging with far reaching impact. Perhaps one of the most relevant topics today is consumer privacy. For example, Microsoft recently published a prospectus entitled "*The Future Computed: Artificial Intelligence and Its Role in Society*." Here the authors frame issues of fair and equitable value that consumers gain from machine interactions with their personal information. Another concern often raised is the reliability and accuracy of algorithms developed from incomplete subsets of data to accurately represent populations. In a focus on the artificial intelligence applications through democratized processes, it is positioned that:

> Could we see a Hippocratic Oath for coders like we have for doctors? That could make sense. We'll all need to learn together and with a strong commitment to broad societal responsibility. Ultimately the question is not only what computers can do. It's what computers should do. Similarly, will the future give birth to a new legal field called 'AI law'? Today AI law feels a lot like privacy law did in 1998.

Also relevant, but currently deeper in the technology is the ambiguity of how machines learn paired with the need for comprehensible explanations of how decisions are made. In development of deep neural networks, machines are essentially writing their own code. Today, the majority of that code cannot be deciphered within a reasonable amount of time by human beings. When machines are producing the decisions and predictions expected, humans are not likely to investigate the code to understand the logic behind them. However, when it matters most, as in times where safety is compromised or lives are lost, accurate explanations in reasonable timeframes need to be required. Such was the case when an individual's Tesla car crashed, killing the driver, while in autopilot mode in 2016 (Banks, Plant and Stanton, 2017). How can we avoid a similar calamity in medical care as a result of incomplete data sets?

While examples exist of rule-based neural network machine learning where pre-labeled samples are not necessary for success (e.g., AlphaGo Zero), the majority of current machine learning requires some form of input of data. This means that not only must processes must be in place to assure the training input, but research must be done now to discover the best possible implementation of the algorithms. Such research projects are emerging worldwide, and include the Computable Evidence Lab, part of the Australian Institute of Health Innovation at Macquarie

University. Their most recent projects are using machine learning and natural language processing to automate systematic reviews, assess and develop robust clinical quality indicators, and better understand the epidemiology of antibiotic resistance.

The recent history of health records, their transformation into electronic health records (EHR), and the ambiguity that currently lives in the data creates significant obstacles in introducing AI effectively into these systems. The advancements in natural language processing coupled with recent developments in artificial intelligence, however, have realized recent successes in interpreting electronic health records. Google was able to accurately predict medical events using only EHR data (Rajkomar et al, 2018). Just over 200,000 patients' records produced more than 46 billion data points to be analyzed. The predictive models were able to accurately predict in-hospital mortality (93% accuracy), 30-day unplanned readmission (75% accuracy), prolonged length of stay (85% accuracy), and patients' final diagnoses (90% accuracy).

These examples demonstrate the potential applicability through research and modeling, with much work to be done by way of implementation. While they certainly show promise in interventions and impact, much work remains in having new ways of providing health care being integrated into our health systems.

DISCUSSION

The rise of personal computing becomes more *personal* every day, and with it the line between e-commerce and e-health is more blurred. Consumers view the digital marketplace, which is continually more tailored to their experience, as the same ground upon which their digital health experience should be marketed and housed. These marketplaces hold vast amounts of data, now both in consumer and consumer health activities. There are fairly few companies that currently hold such a large stake of the world's data. AI and big data are opening the opportunity for health care and health care delivery for new corporate alignment and opportunities (Reilly, 2018). Concurrently, multiple sources are reporting on data as one of the most traded commodities in the world and increasing every year and for the first time overcoming materials, manufactured goods, and services (Economist, 2017). As we are truly heading into existing in a "knowledge economy" and a world based on analytics, and data and automation, there are profound implications for governments and society that are at hand.

The technologies and applications discussed herein have been mostly conceptual in their widespread application in medical care until now despite technologies like natural language processing have been researched for more than 35 years (Shortliffe et al, 1973). The emergence, however, of entirely new and enabling technologies (digital health information, telehealth applications, broad band digital transmission, and cloud computing, etc.) have created the resources and capabilities to revolutionize how health care is delivered. The implications of providing medical knowledge and high quality, low cost care to rural communities and the developing regions of the world have the potential to have profound implications in global health. The framework for such revolution is albeit undefined and rather unstructured, as most technologies were developed without integrative or interdependent considerations. Extensive work is needed to enhance the reliability and safety of aided-diagnosis and medical decision-making, for example. With the right support and design, these applications can overcome barriers to current health care problems like access to care, translation of evidence and knowledge into practice, improving health literacy and overcoming communication barriers.

The aspects of human to human engagement in health will challenge existing models of adoption, regulatory oversight, transparency, and equity. The GDHP represents a forum for addressing the evolutionary roles of these technologies and the considerations of human values through the digital lens of health information toward the goal of providing a greater good for all of humanity.

# # #

REFERENCES

Altae-Tran, Han, Bharath Ramsundar, Aneesh S. Pappu, and Vijay Pande. "Low data drug discovery with one-shot learning." ACS central science 3, no. 4 (2017): 283-293.

"Artificial intelligence for health and healthcare." *JASON White Paper* (2017).

Banks, Victoria A., Katherine L. Plant, and Neville A. Stanton. "Driver error or designer error: Using the Perceptual Cycle Model to explore the circumstances surrounding the fatal Tesla crash on 7th May 2016." Safety Science (2017).

Brown, Peter, John Cocke, S. Della Pietra, V. Della Pietra, Frederick Jelinek, Robert Mercer, and Paul Roossin. "A statistical approach to language translation." In *Proceedings of the 12th conference on Computational linguistics-Volume 1*, pp. 71-76. Association for Computational Linguistics, 1988.

"CAMELYON17 - GrandChallenge.org." CAMELYON17 - Home. January 2017. Accessed February 04, 2018. https://camelyon17.grand-challenge.org/.

"Computable evidence lab | AIHI - Australian Institute of Health Innovation." Macquarie University Research, Higher Degree Research & Innovation. Accessed February 04, 2018. http://aihi.mq.edu.au/content/computable-evidence-lab.

"The Connectivity Map." Clue.io. Accessed February 06, 2018. https://clue.io/.

The Future Computed: Artificial Intelligence and Its Role in Society. Redmond, WA: Microsoft, 2018. January 17, 2018. Accessed January 31, 2018.

Gao, Jim, and Ratnesh Jamidar. "Machine learning applications for data center optimization." *Google White Paper* (2014).

Garvin, Jennifer Hornung, Youngjun Kim, Glenn Temple Gobbel, Michael E. Matheny, Andrew Redd, Bruce E. Bray, Paul Heidenreich et al. "Automating Quality Measures for Heart Failure Using Natural Language Processing: A Descriptive Study in the Department of Veterans Affairs." *JMIR medical informatics* 6, no. 1 (2018): e5.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press.

Herper, Matthew. "MD Anderson Benches IBM Watson In Setback For Artificial Intelligence In Medicine." Forbes. February 24, 2017. Accessed January 31, 2018. https://www.forbes.com/sites/matthewherper/2017/02/19/md-anderson-benches-ibm-watson-in-setback-for-artificial-intelligence-in-medicine/#1462c0d33774.

Hickman, Rodney. "Front-End Speech What are the Value-Adds?" Front-End Speech What are the Value-Adds? - KLAS Report. March 2015. Accessed January 31, 2018. https://klasresearch.com/report/front-end-speech/1027.

"Machine Learning Report." Report: Introduction to Machine Learning. 2017. Accessed January 31, 2018. https://orionhealth.com/nz/knowledge-hub/reports/machine-learning-report/.

McCulloch, Warren S., and Walter Pitts. "A logical calculus of the ideas immanent in nervous activity." *The bulletin of mathematical biophysics* 5, no. 4 (1943): 115-133.

Rajkomar, Alvin, Eyal Oren, Kai Chen, Andrew M. Dai, Nissan Hajaj, Peter J. Liu, Xiaobing Liu et al. "Scalable and accurate deep learning for electronic health records." *arXiv preprint arXiv:1801.07860* (2018).

Reilly, Mark. "UnitedHealth shares dive as Amazon, Berkshire Hathaway, JP Morgan Chase team up on health care." Bizjournals.com. January 30, 2018. Accessed February 04, 2018.https://www.bizjournals.com/twincities/news/2018/01/30/unitedhealth-shares-plummet-as-amazon-berkshire.html.

Ross, Casey, and Ike Swetlitz. "IBM pitched Watson as a revolution in cancer care. It's nowhere close." STAT. September 05, 2017. Accessed January 31, 2018. https://www.statnews.com/2017/09/05/watson-ibm-cancer/.

Shortliffe, Edward H., Stanton G. Axline, Bruce G. Buchanan, Thomas C. Merigan, and Stanley N. Cohen. "An artificial intelligence program to advise physicians regarding antimicrobial therapy." *Computers and Biomedical Research* 6, no. 6 (1973): 544-560.

Silver, David, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert et al. "Mastering the game of go without human knowledge." *Nature* 550, no. 7676 (2017): 354.

Turing, Alan M. "Computing machinery and intelligence." In *Parsing the Turing Test*, pp. 23-65. Springer, Dordrecht, 2009.

"Up And Up: Healthcare AI Startups See Record Deals." Research Briefs. August 22, 2017. Accessed February 02, 2018. https://www.cbinsights.com/research/artificial-intelligence-healthcare-deals-funding-investors/.

Vitalari, Nicholas P., and Haydn Shaughnessy. *The elastic enterprise: the new manifesto for business revolution*. Longboat Key, FL: Telemachus Press, 2012.

"The world's most valuable resource is no longer oil, but data." The Economist. May 06, 2017. Accessed February 04, 2018. https://www.economist.com/news/leaders/21721656-data-economy-demands-new-approach-antitrust-rules-worlds-most-valuable-resource.

"Your.MD." Your.MD - Home. Accessed February 04, 2018. https://www.your.md/.

APPENDIX

A Brief History of Artificial Intelligence

1805

Adrien-Marie Legendre describes the “method des moindres carres” known in English as the least squares method – used widely in data fitting.

1929

Makoto Nishimura designs Gakutensoku, Japanese for "learning from the laws of nature," the first robot built in Japan. It could change its facial expression and move its head and hands via an air pressure mechanism.

1943

Warren McCulloch & Walter Pitts publish "A Logical Calculus of the Ideas Immanent in Nervous Activity," laying foundations for neural networks.

1946

John Von Neumann Electronic Numerical Integrator & Computer (ENIAC), prompted Alan Turing to design his Turing Test.

1950

Alan Turing published "Computing Machinery and Intelligence," also known as The Turing Test.

1952

Arthur Samuel develops the world’s first self-learning program Checkers learning program.

1955

Herbert Simon and Allen Newell develop the Logic Theorist, the first artificial intelligence program, which eventually would prove 38 of the first 52 theorems in Whitehead and Russell's *Principia Mathematica*.

1958

John McCarthy develops programming language Lisp which becomes the most popular programming language used in artificial intelligence research.

1963

Ivan Sutherland's MIT dissertation on Sketchpad introduced the idea of interactive graphics into computing.

Edward A. Feigenbaum & Julian Feldman published Computers and Thought, the first collection of articles about artificial intelligence.

1964
Daniel Bobrow completes his MIT PhD dissertation titled “Natural Language Input for a Computer Problem Solving System” and develops STUDENT, a natural language understanding computer program.

1967
Richard Greenblatt at MIT builds a knowledge-based chess-playing program, MacHack, that is good enough to achieve a class-C rating in tournament play.

1969
First International Joint Conference on Artificial Intelligence (IJCAI) held in Washington, D.C.

1973
The Assembly Robotics group at Edinburgh University builds Freddy, the Famous Scottish Robot, capable of using vision to locate and assemble models.

1978
Herb Simon wins the Nobel Prize in Economics for his theory of bounded rationality, one of the cornerstones of AI known as "satisficing".

1979
The Stanford Cart, built by Hans Moravec, becomes the first computer-controlled, autonomous vehicle when it successfully traverses a chair-filled room and circumnavigates the Stanford AI Lab.

1980’s
Lisp Machines developed and marketed.

James Allen invents the Interval Calculus, the first widely used formalization of temporal events.

Mid 80's
Neural Networks become widely used with the Backpropagation algorithm (first described by Werbos in 1974).

1986
First driverless car, a Mercedes-Benz van equipped with cameras and sensors, built at Bundeswehr University in Munich under the direction of Ernst Dickmanns, drives up to 55 mph on empty streets.

1988
Members of the IBM T.J. Watson Research Center publish “A statistical approach to language translation,” heralding the shift from rule-based to probabilistic methods of machine translation, successfully translating between English and French based on 2.2 million pairs of sentences, mostly from the bilingual proceedings of the Canadian parliament.

1997

The Deep Blue chess program beats the current world chess champion, Garry Kasparov, in a widely followed match and rematch.

NASA's pathfinder mission makes a successful landing and the first autonomous robotics system, Sojourner, was deployed on the surface of Mars.

2009

Google starts developing, in secret, a driverless car. In 2014, it became the first to pass, in Nevada, a U.S. state self-driving test.

2010

Kaggle Website is launched, that serves as a platform for machine learning competitions.

2011

A convolutional neural network wins the German Traffic Sign Recognition competition with 99.46% accuracy (vs. humans at 99.22%).

IBM Watson defeats two human challengers at Jeopardy! It used a combination of machine learning, natural language processing and information retrieval techniques.

2016

Google DeepMind's AlphaGo defeats Go champion Lee Sedol.

Machine learning widespread Evans Data's Big Data and Advanced Analytics Survey found that more than one-third of developers say they're using machine learning technologies in their big data and advanced analytics projects.