

REPORT OF THE COUNCIL ON LONG RANGE PLANNING AND DEVELOPMENT

The following report was presented by Glenn Loomis, MD, Chair.

1. A PRIMER ON ARTIFICIAL AND AUGMENTED INTELLIGENCE

Informational report; no reference committee hearing.

HOUSE ACTION: FILED

Last year, the Council on Long Range Planning and Development (CLRPD) created the educational module, Health Care Trends: Scientific Innovation,¹ which accelerated its interest in artificial and augmented intelligence (AI), and prompted a series of discussions on these topics and their influences on the practice of medicine. Due to the complexity of the field, the Council developed this primer, which provides a history, definitions and components, and the status of AI in health care. Additionally, CLRPD postulated ways the field may progress, including the identification of opportunities and challenges for physicians. The Council feels it essential to provide a high-level look at this emerging issue that could dramatically affect medicine.

HISTORY OF AI

The most influential ideas underpinning computer science came from Alan Turing in 1950, who proposed a formal model of computing. Turing's classic essay, *Computing Machinery and Intelligence*,² imagines the possibility of computers created for simulating intelligence and explores many of the components now associated with artificial intelligence, including how intelligence might be tested, and how machines might automatically learn. Though these ideas inspired AI, Turing did not have access to the computing resources needed to translate his ideas into action.

In 1956, the field of AI came to the forefront with the Dartmouth Summer Research Project on Artificial Intelligence. The goal was to investigate ways in which machines could be made to simulate aspects of intelligence—the essential idea that has continued to drive the field forward. Subsequently, experts in the field of computer science research pioneered the foray into heuristic search—a method that produces a solution in a reasonable timeframe that is sufficient for solving a given problem. In the area of computer vision, early work in character recognition laid the basis for more complex applications such as face recognition. By the late sixties, work had also begun on natural language processing (NLP).

In the nineties, technological progress made the task of building systems driven by real-world data more feasible. Cheaper and more reliable hardware for sensing and actuation made robots easier to build. Further, the Internet's capacity for gathering large amounts of data, and the availability of computing power and storage to process those data enabled statistical techniques that, by design, derive solutions from data. These developments have allowed AI to emerge in the past two decades as a profound influence on our daily lives.³

DEFINITIONS AND COMPONENTS OF AI

The concepts of AI and machine learning have quickly become attractive to health care organizations; however, the related terminologies are not well understood. While many in the health care industry foresee their technological goals hovering just over the horizon, plotting a course to get there can be a difficult proposition, especially when the landscape is clouded by marketing hyperbole, confusing vocabulary, technical terminology, and as-yet-undeliverable promises of truly automated insights.

Algorithms are a sequence of instructions used to solve a problem. Developed by programmers to instruct computers in new tasks, algorithms are the building blocks of the advanced digital world. Computer algorithms organize enormous amounts of data into information and services, based on certain instructions and rules.

Artificial Intelligence is the ability of a computer to complete tasks in a manner typically associated with a rational human being—a quality that enables an entity to function appropriately and with foresight in its environment. True AI is widely regarded as a program or algorithm that can beat the Turing Test, which states that an artificial intelligence must be able to exhibit intelligent behavior that is indistinguishable from that of a human.

Augmented Intelligence is an alternative conceptualization that focuses on AI's assistive role, emphasizing the fact that its design enhances human intelligence rather than replaces it.

Machine Learning is a part of the discipline of artificial intelligence and refers to constructing algorithms that can make accurate predictions about future outcomes. Machine learning can be supervised or unsupervised. In supervised learning, algorithms are presented with “training data” that contain examples with their desired conclusions, such as pathology slides that contain cancerous cells as well as slides that do not. Unsupervised learning does not typically leverage labeled training data. Instead, algorithms are tasked with identifying patterns in data sets on their own by defining signals and potential abnormalities based on the frequency or clustering of certain data.

Deep Learning is a subset of machine learning that employs artificial neural networks (ANNs) and algorithms structured to mimic biological brains with neurons and synapses. ANNs are often constructed in layers, each of which performs a slightly different function that contributes to the result. Deep learning is the study of how these layers interact and the practice of applying these principles to data.

Cognitive Computing, a term coined by IBM, is often used interchangeably with machine learning and artificial intelligence. However, cognitive computing systems do not necessarily aspire to imitate intelligent human behavior, but instead to supplement human decision-making power by identifying potentially useful insights with a high degree of certainty. Clinical decision support and augmented intelligence come to mind when considering this definition.

Natural Language Processing (NLP) forms the foundation for many cognitive computing exercises. The ingestion of source materials, such as medical literature, clinical notes, or audio dictation records requires a computer to understand what is written, spoken or otherwise being communicated. One commonly used application of NLP is optical character recognition (OCR) technology that can turn static text, such as a PDF of a lab report or a scan of a handwritten clinical note, into machine readable data. Once the data are in a workable format, the algorithm parses the meaning of each element to complete a task such as translating into a different language, querying a database, summarizing information or supplying a response to a conversation partner. In the health care field, where acronyms and abbreviations are common, accurately parsing through this “incomplete” data can be challenging.

On a basic level, classical computer programming takes *rules* and *data* as inputs, and generates an output or answer. Conversely, machine learning algorithms take *data* and *answers* as inputs, and generate *rules* or *insights* as an output. For example, a computer may be given two sets of MRI images: one set that clearly shows a variety of brain tumors, and one that does not. By breaking down these images into machine-readable patterns, the computer can understand which patterns are likely to indicate a brain tumor and which represent healthy patients. When fed a new batch of images that may or may not contain tumors, the computer should be able to use that initial reference data to identify patterns that are similar to known positive diagnoses. Every time it makes an incorrect diagnosis, validated by a human clinician, it “learns” to adjust its criteria a little bit more by using the previous experience to inform its future decision-making. With enough training, it can become accurate enough to present reliable results to the user.

Humans complete these types of tasks almost without thought every moment of every day, but few algorithms are sophisticated enough to effectively mimic our natural capacity to process external input, extrapolate unspoken information from a query, consider complex ethical issues, use logic and reason to make a decision, and predict the likely outcomes of each action before they occur. When comparing the common definition of AI as the capability of a machine to imitate intelligent human behavior with the Turing Test challenge of creating an algorithm that performs a task indistinguishably from a human counterpart, it becomes clear that machines are still in the process of evolving. However, there are a few examples of use cases in health care that are coming closer to realizing the Turing Test.

STATUS OF HEALTH CARE AI

Some of the most promising use cases for health care AI tools include predictive analytics, precision medicine, and clinical decision support. Development in all of these areas is already well underway. The private sector has acknowledged these opportunities, and investments in AI have grown over the past several years.⁴ A recent report from Markets and Markets pins the health care AI sector at nearly \$8 billion in 2022, accelerating at a compound annual growth rate of 52.68 percent over the forecast period.⁵

In 2011, IBM got an early start in the health care AI space by using Watson’s NLP and cognitive computing abilities to train in clinical decision support at some of the top medical institutions in the country. IBM has also committed extensive resources, such as its \$2.6 billion acquisition of Truven, to imaging analytics, genomics, pharmaceuticals,

and population health management.⁶ Their efforts are not without roadblocks—a multiyear project to apply IBM Watson to cancer diagnostics with MD Anderson ended in failure.⁷ Other industry leaders, Google and Microsoft, are ramping up their efforts to apply advanced machine learning algorithms to the mysteries of human biology. Microsoft is tackling genomics, cancer, myopia and blindness, transplants, and imaging analytics,⁸ while Google recently published research on the role of machine learning in pathology and breast cancer,⁹ and diabetic retinopathy.¹⁰ Additionally, Google is the first of the titans to establish a formal program, Launchpad Studio, for working with startups specific to the industry, such as Augmedix, BrainQ, Byteflies, and Cytovale.¹¹

Currently, machine learning has started to prove its value in the realm of pattern recognition, NLP, and deep learning. At the Stanford University School of Medicine, a machine-learning algorithm out-performed pathologists at predicting patient survival times for two types of lung cancer.¹² In the United Kingdom, a NLP tool applied to free-text peer assessments of physician performance, derived by human raters, agreed with the content of the documents 98 percent of the time.¹³ At Indiana University-Purdue University Indianapolis, machine learning correctly predicted relapse rates for a type of leukemia 90 percent of the time. It identified patients who would experience remission with 100 percent accuracy.¹⁴ Engineers at Boston University are working with Brigham and Women’s Hospital, and Boston Medical center to manage heart diseases and diabetes using algorithms that have the ability to predict hospitalizations up to a year in advance with 82 percent accuracy.¹⁵ However, current algorithms do not result in autonomous decisions. Instead, they play an assistive role to augment human intelligence rather than replace it.

FUTURE OF AI IN MEDICINE

What does the future of AI hold in medicine? AI technology could change the world for the better by making care delivery safer, improving diagnostic accuracy, increasing physician productivity and scale, or contributing to applications that improve quality of life. As the technology of AI continues to develop, physicians and medical associations must ensure that AI-enabled systems are governable; are open, transparent, and understandable; can work effectively with people; are included in medical education for students and practicing physicians; and remain consistent with human and medical ethics. Physician involvement with the evolution of this active field may help them to chart a better and wiser path forward for themselves, their patients, and the health care system.

Opportunities and challenges of AI in health care are equally profound for physicians:

Opportunities

- Office and hospital automation – patient scheduling, order entry, chat bots, voice recognition, etc.
- Data mining to surface the right data at the right time, and improve EHRs
- Diagnosis – analyze all the known data about the patient and produce insights
- Treatment – analyze the diagnosis and all other known data and produce best practice treatments, perhaps even comparing to “patients like me” data
- Additional time for physicians to spend with patients to focus on their health
- Improve patient experience, and aid behavioral change and treatment compliance
- Medical education – personal assistant for students and residents to surface information (less memorization), automated continuous assessment of competencies, and coaching

Challenges

- Data structure, integrity and security
- Technological mistrust – transparency is key
- Demonstrate that AI can reduce costs, deliver the quadruple aim, support the patient-physician relationship, and/or alleviate administrative burden
- Implement and integrate AI into clinical practices and patient care
- Uncertain long term employment outlook for health care professionals
- Susceptibility to training bias, malfeasance, and other possible technical problems
- Questions as to who will benefit and who may lose—what is best for an individual is not always best for public health, especially when limited resources are available

Additionally, AI opportunities and challenges lead to questions physicians will need to confront:

- What evidence is needed to demonstrate value, utility, and trust?
- How does AI intersect with other emerging health care capabilities, such as genomic medicine?
- How will regulatory bodies and professional organizations provide proper oversight for AI benefits and risks, and communicate these to the public?
- How can public and systemic expectations be managed, and concerns allayed?
- What education and training will health care professionals need to acquire in order to understand how AI solutions might help them, and their patients in clinical settings?
- What can health systems considering AI opportunities do now to maximize their chances of success for gaining efficiencies, improving care, and integrating into clinical workflows?
- How will risk be allocated, given the “black box” nature of AI systems?
- How will legal, policy, and regulatory implications, including standards for professional services, intellectual property rights, and FDA oversight be monitored and addressed?

Beyond the potential to dramatically affect the economy and society in the near future, AI has moved to the forefront of many policy debates around the world. These debates range from the governance of AI, such as ensuring accountability of algorithmic decisions, to mitigating the impact of AI on employment. Clear challenges must be addressed to support AI’s future in medicine. Therefore, it is up to all stakeholders, be they health care professionals, medical associations, policymakers, businesses, the technology industry, or civil society to ensure that AI’s impact is a positive one by proactively tackling the challenges, while ensuring the opportunities remain available.

REFERENCES

1. AMA-CLRPD. *Health Care Trends: Scientific Innovation*. <https://cme.ama-assn.org/Activity/5617091/Detail.aspx>.
2. Turing, AM. *Computing Machinery and Intelligence*. <http://phil415.pbworks.com/f/TuringComputing.pdf>.
3. Stanford. *One Hundred Year Study on Artificial Intelligence, 2016 Report*. <https://ai100.stanford.edu/2016-report>.
4. CB Insights. *The Race for AI: Google, Baidu, Intel, Apple in a Rush to Grab Artificial Intelligence Startups*. <https://www.cbinsights.com/research/top-acquirers-ai-startups-ma-timeline/>.
5. MarketsandMarkets. *Artificial Intelligence in Healthcare Market by Offering (Hardware, Software and Services), Technology (Deep Learning, Querying Method, NLP, and Context Aware Processing), Application, End-User Industry, and Geography – Global Forecast to 2022*. https://www.marketsandmarkets.com/Market-Reports/artificial-intelligence-healthcare-market-54679303.html?gclid=CjwKCAiAr_TQBRB5EiwAC_QCq3mVjD0-a0Lo3XVAu-TwQC4n8H5xUriz6q9BSmKYT2J_WfPwNLA0VBoC6vQQAuV_D_BwE.
6. IBM. *IBM Watson Health Closes Acquisition of Truven Health Analytics*. <https://truvenhealth.com/media-room/press-releases/detail/prid/185/ibm-watson-health-announces-plans-to-acquire-truven-health-analytics>.
7. ARS Technica. *IBM’s Watson Proves Useful at Fighting Cancer-Except in Texas*. <https://arstechnica.com/science/2017/02/ibms-watson-proves-useful-at-fighting-cancer-except-in-texas/>.
8. Xconomy. *Microsoft’s Strategy for Finding What’s Next in Healthcare AI*. <https://www.xconomy.com/seattle/2017/11/08/microsofts-strategy-for-finding-whats-next-in-healthcare-a-i/>.
9. Google Inc. *Detecting Cancer Metastases on Gigapixel Pathology Images*. <https://arxiv.org/abs/1703.02442>.
10. GulshanV, et al. *Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs*. <https://jamanetwork.com/journals/jama/fullarticle/2588763>.
11. HealthIT News. *Google powers up AI, machine learning accelerator for healthcare*. <http://www.healthcareitnews.com/news/google-powers-ai-machine-learning-accelerator-healthcare>.
12. Yu KH, et al. *Predicting non-small cell lung cancer prognosis by fully automated microscopic pathology image features*. <https://www.nature.com/articles/ncomms12474>.
13. Gibbons C, et al. *Supervised Machine Learning Algorithms Can Classify Open-Text Feedback of Doctor Performance With Human-Level Accuracy*. <http://www.jmir.org/2017/3/e65/>.
14. Rajwa B, et al. *Automated Assessment of Disease Progression in Acute Myeloid Leukemia by Probabilistic Analysis of Flow Cytometry Data*. <https://www.ncbi.nlm.nih.gov/pubmed/27416585>.
15. Dai W, et al. *Prediction of hospitalizations due to heart diseases by supervised learning methods*. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4314395/pdf/nihms643126.pdf>.