VIEWPOINT

On the Prospects for a (Deep) Learning Health Care System

C. David Naylor, MD, DPhil
Department of Medicine, University of Toronto, Toronto, Ontario, Canada.

Viewpoint and Editorial

Supplemental content

Corresponding Author: C. David Naylor, MD, DPhil, Department of Medicine, University of Toronto, 4 Shorncliffe Ave, Toronto, ON M4V 1T1, Canada (david.naylor@utoronto.ca).

In 1976, Maxmen[1] predicted that artificial intelligence (AI) in the 21st century would usher in "the post-physician era," with health care provided by paramedics and computers. Today, the mass extinction of physicians remains unlikely. However, as outlined by Hinton[2] in a related Viewpoint, the emergence of a radically different approach to AI, called deep learning, has the potential to effect major changes in clinical medicine and health care delivery. This Viewpoint reviews some of the factors driving wide adoption of deep learning and other forms of machine learning in the health ecosystem.

Emergence of Deep Learning

Forms of machine learning (eg, linear and logistic regression models) have been applied for decades in basic and medical research. The winning conditions for more advanced forms of machine learning developed in the late 1990s as computational devices became ubiquitous and more mobile, interconnected, and powerful. Mass data compilation and retrieval were facilitated by remote storage and computing, the phenomenon later labeled as "the cloud." The health realm in particular saw exponential growth in the volume of information concerning patients and their journey through the world's health care systems from large administrative data sets, genomics and related disciplines, advanced imaging techniques, handheld devices that enabled clinicians to create instant digital records, and the broad public uptake of "wearables." Thus, in a range of disciplines, not least medicine and health care, data mining substantially expanded.

As Shortliffe et al[3] observed in 2009, advances in symbolic or rule-based AI and statistically derived forms of machine learning were often subsumed in these technologic shifts, without much public acknowledgment. Examples included accelerated gene sequencing and automated searches of massive biodata banks, computerized enhancement of medical images from a wide array of devices, knowledge management systems to concatenate clinical records and facilitate collaborative care, implantable cardiac defibrillators, sophisticated intensive care unit monitoring alarms, and the endless variety of automated processes for procurement, scheduling, and drug ordering in thousands of clinics and hospitals worldwide.

Artificial neural networks, a model for AI patterned loosely on the mammalian visual cortex, made few inroads during much of this period. However, starting in approximately 2006, technical advances and much faster hardware made it feasible to train neural networks with many layers of feature detectors on large data sets. The term *deep learning* was adopted to differentiate this new generation of neural net technology from its progenitors. In 2012, deep learning developed at the University of Toronto was used to significantly improve speech recognition on Android devices and the prediction of drug activity in a Merck competition. It also substantially outperformed other approaches to computer vision in a public competition.

The technology quickly found wide commercial use. Early adopters included Google, Facebook, Microsoft, Apple, and Amazon, with the result that deep learning became globally ubiquitous almost overnight.

Deep learning had intuitive appeal for health-related applications, given its demonstrable strengths in intricate pattern recognition and predictive model building from big high-dimensional data sets. These analytic capabilities have already proven useful for basic and applied researchers, ranging across health disciplines. Thus far, clinical application of deep learning has been most rapid in image-intensive fields such as radiology, radiotherapy, pathology, ophthalmology, dermatology, and image-guided surgery. In many cases, interpretation of images by deep learning systems has outperformed that by individual clinicians when measured against a consensus of expert readers or gold standards such as pathologic findings. Clinically relevant applications have widened beyond image processing to include risk stratification for a broad range of patient populations (eBox in the Supplement), and health care organizations are capitalizing on deep learning and other machine-learning tools to improve logistics, quality management, and financial oversight.

Factors Driving Adoption of AI and Deep Learning

Seven factors appear to be driving the rapid uptake of AI and deep learning.

1. *Digital imaging in all its forms is becoming more powerful and more integral to medicine and health care*. Unlike deep learning, expert human interpretation fails to capitalize on all the patterns, or "regularities," that can be extracted from very large data sets and used for interpretation of still and moving images. Deep learning and related machine-learning methods can also learn from massively greater numbers of images than any human expert, continue learning and adapting over time, mitigate interobserver variability, and facilitate better decision making and more effective image-guided therapy.

2. *Digitization of health-related records is accelerating, as is sharing of high-quality fully labeled and specialized data sets*. These data sources provide new opportunities for application of data science and refinement of algorithms. Hoarding and commercialization of health data by for-profit corporations remains a potential impediment. However, collaborative groups of basic and clinical researchers have shown impressive leadership in sharing and curating data sets, as well as convening contests under open-access rules

to provide incentives for the development, comparison, and sharing of different machine-learning models.

3. *Deep learning is highly adaptable for integrative analysis of heterogeneous data sets assembled from diverse sources*. Combi[4] described heterogeneous data as arising from concatenation of "demographic data, temporal clinical and health data, biomedical signals and images, genetic data, biomolecular data, clinical pathway data, [and] social network data, just to mention some wide categories of data and without making any claim to completeness." Those composite records, in turn, can be aggregated at multiple levels: across groups of patients for scientific research and technology development with or without enrichment by additional measures; across teams of health care professionals for management of safety, quality, and productivity; across institutions for comparative performance profiling; and across health care systems to inform public policy making.

4. *Deep learning has enormous capacity to inform the process of discovery in health research and to facilitate hypothesis generation by identifying novel associations*. Established and start-up companies are using deep learning to select or design novel molecules for testing as pharmaceuticals or biologics, with in silico exploration preceding in vitro examination and in vivo experimentation. Researchers across disciplines have also found unexpected clusters within data sets by comparing the intensity of activation of feature detectors in the hidden layers of deep neural nets. As always, however, basic and clinical experimentation remains essential to establish causation and causal pathways.

5. *Deep learning shows promise for streamlining routine work by health care professionals and empowering patients, thereby promoting a safer, more humane, and participatory paradigm for health care.* Different sources offer varying estimates of the amount of time wasted by health care professionals on tasks amenable to some automation (eg, high-quality image screening) that could then be rededicated to more or better care. A growing number of research studies also suggest specific possibilities for reduction in errors and improved work flow in the clinical setting with appropriate deployment of AI.

Combined with wearables, remote monitoring, and digital consultations, deep learning and other machine-learning techniques can bypass the time-honored model of intermittent data collection and interpretation at the clinical encounter. These advances may promote more effective and informed self-care by patients and families.[5] In the longer term, deep learning can relate those personalized features to the clinical course of similar patients, using data from millions of patient records containing billions of medical events. Thus, while concerns are understandably raised that automation could dehumanize clinical care, these advances could provide professionals and patients alike with vastly better and more specific information, and, as Fogel and Kvedar[6] argue, give physicians more time "to focus on the tasks that are uniquely human: building relationships, exercising empathy, and using human judgment to guide and advise."

6. *Deep learning is diffusing rapidly through a combination of open-source and proprietary programs*. Technology giants are making massive investments in the development of software libraries for deep learning, some of which are open sourced. These huge enterprises, as well as start-ups, are applying deep learning tools to health care all over the world. Moreover, many academic and nonprofit teams are publishing and sharing algorithms freely, and local development is now widespread.

Regulators are rightly trying to ensure that appropriate evidence supports claims for various commercial software packages, a particularly urgent concern for mobile applications targeting consumers. Many publications report on pilot projects with uncertain feasibility for wide clinical adoption. For some interventions (eg, AI-guided robotic-assisted surgery), evidence will be needed from randomized trials. For other applications, the standard approach used for evaluating diagnostic tests seems more appropriate. However, unlike a standardized diagnostic test or drug, the performance of deep learning and other machine-learning methods improves with exposure to larger or more relevant data sets, or with easily made modifications to the architecture of the models or training procedures. Regulators and technology assessors will need to distinguish issues inherent in decision-support algorithms from those attributable to misuse by clinical decision makers. Procurement agencies and health care administrators will need to be uncharacteristically nimble to keep up.

7. *None of these developments depend on improvements in the basic technology of deep learning*. As Hinton[2] observes, ongoing technical refinements are likely to make deep learning more efficient and effective in the years ahead. Regardless of technical refinements, however, deep learning performance will improve as data sets become larger, better linked and labeled, and more detailed.

Conclusions

While deep learning at present has some unique advantages as an analytic and modeling tool, it also exemplifies a broader trend: the convergence of health and data sciences. Barriers to adoption will rightly be more rigid in health care than in many other fields in which software programs relying on deep learning and other forms of machine learning are used daily by billions of people. However, pressure to deploy deep learning and a range of tools derived from modern data science will be relentless, given the extraordinarily rich information now available to characterize and follow vast numbers of patients, the ongoing challenges of making sense of the complexity of human biology and health care systems, and the potential for smart information technology to support tomorrow's clinicians in the provision of safe, effective, efficient, and humanistic care.

ARTICLE INFORMATION

Published Online: August 30, 2018. doi:10.1001/jama.2018.11103

Conflict of Interest Disclosures: The author has completed and submitted the ICMJE Form for Disclosure of Potential Conflicts of Interest and none were reported.

Additional Contributions: I am deeply indebted to Geoffrey Hinton, PhD, for discussions that shaped this Viewpoint and for sharing insights into advances in AI that seemed unthinkable when we last collaborated more than 20 years ago. This Viewpoint is dedicated to the memory of Jack Ven Tu, MD, PhD (1965-2018).

REFERENCES

1. Maxmen JS. *The Post-Physician Era: Medicine in the 21st Century*. New York, NY: Wiley; 1976.

2. Hinton G. Deep learning: a technology with the potential to transform health care [published online August 30, 2018]. *JAMA*. doi:10.1001/jama.2018.11100

3. Patel VL, Shortliffe EH, Stefanelli M, et al. The coming of age of artificial intelligence in medicine. *Artif Intell Med*. 2009;46(1):5-17.

4. Combi C. Editorial from the new editor-in-chief: artificial intelligence in medicine and the forthcoming challenges. *Artif Intell Med*. 2017;76:37-39.

5. Topol E. *The Patient Will See You Now: The Future of Medicine Is in Your Hands*. New York, NY: Basic Books; 2015.

6. Fogel AL, Kvedar JC. Artificial intelligence powers digital medicine. *NPJ Digital Medicine*. 2018; 1:5. doi:10.1038/s41746-017-0012-2