Artificial Intelligence and the Practice of Radiology: An Alternative View

Robert Schier, MD

INTRODUCTION

Massive transformations often begin with deceptive slowness. The first hand-cranked digital computer in 1672 could add and subtract. The first electronic digital computer in 1946 was also able to multiply and divide. In 1956, researchers in the newly named field of artificial intelligence (AI) predicted that computers as intelligent as humans were just around the corner. But it took until 1997 for a computer to beat the best human chess player and until 2016 for a computer to win at the vastly more complicated game of Go. Within just the past few years, computers have begun to equal or exceed human abilities in an increasing range of tasks. Computer intelligence is now beginning to impact many areas in the practice of medicine; the recent rapid advances in computer interpretation of images suggest that radiology will be affected sooner than most.

BEING INTELLIGENT AND SOLVING PROBLEMS

Will machine intelligence eventually surpass human intelligence to the same extent that machine power surpassed animal power? More important, will computers eventually take over most jobs humans now do just as machines took over jobs that horses, donkeys, and oxen once did? The answer to both questions seems increasingly likely to be yes. Three years ago, Stephen Hawking, Elon Musk, and a number of researchers in computer intelligence signed an open letter [1] warning that the advent of intelligent computers will carry with it "potential pitfalls"—Musk himself described one of those pitfalls as computers becoming an "existential threat" to humanity [2]. This latter point is furiously debated. The underlying assumption—that computers will become intelligent enough to warrant such concern—is not.

Being intelligent has nothing to do with "being human"—possessing thoughts, feelings, desires, and emotions. *Intelligence* refers to the ability to solve problems. In limited applications, machines have been more intelligent than humans for years: gear-driven adding machines can solve equations faster and more accurately than any person. But can computers solve the more difficult types of problems that we generally assume require a different and more *real* kind of intelligence—problems that they have not been told how to solve, or that do not have a known solution? It is becoming ever clearer that they can. Computer brains made of silicon and circuits are turning out to be comparable if not better than human brains made of carbon and neurons in solving many types of problems. Though computers do not understand the concept of time, they can plan for the future. Though computers are not conscious, they can be creative, cooperative, and curious [3,4].

Problems come in two types: those for which the method of solving them is known, and those for which it is not. Different methods are used to solve each type: deductive reasoning is used for the first, and inductive reasoning is used for the second. Deductive reasoning uses a known general rule to solve a specific problem of the type described by that rule. Inductive reasoning analyzes many specific events to discover the underlying general rule describing all events of that type.

People have an easier time with deductive reasoning. Any high school physics student can use Newton's Law of Universal Gravitation to determine the gravitational attraction between the moon and the earth. But it required the genius of Isaac Newton, using inductive reasoning, to observe the seemingly unrelated motions of the earth, moon, planets, and other moving objects to (1) realize that those motions were not in fact unrelated and (2) discover the law governing the motions of them all.

HOW COMPUTERS LEARN

Computers have an easier time with deductive reasoning, too. That is what computers were built to do in the first place. Computers learn deductive reasoning by being "taught"—a human writes a program specifying the precise steps needed to solve a particular type of problem.

1546-1440/18/$36.00 ■ https://doi.org/10.1016/j.jacr.2018.03.046

Once a computer is programmed in this way, it can solve any problem of that kind. Computers are good at this kind of problem solving; current computers can solve in 1 second problems too complex for a human to solve given the age of the universe.

Though this approach has proven hugely effective, teaching computers in this way runs into difficulties when (1) the problem to be solved is so complicated that it takes an exorbitant amount of time and resources to write the program, (2) the exact steps needed to solve the problem are not known in enough detail to specify them, or (3) it is not known with certainty whether the problem even has a solution.

The response to this difficulty has been the development of programs with the ability to learn inductively. Inductive programs, to a greater or lesser degree, are able to learn on their own.

There are two different approaches to creating computers capable of inductive reasoning. The first is to write programs with the ability to look for patterns in data. This approach is widely used in data mining and mathematical analysis. These programs are taught how to look for patterns.

The second approach is to write computer programs not only with the ability to look for patterns in data, but with the ability to change themselves to better find those patterns. Once started, these types of programs may have little or no need for a human programmer—they will teach themselves how to look for patterns.

The advent of inductive programs, along with increased processing speed, is the reason that computers have recently become so good at interpreting images. The development of programs capable of inductive reasoning has given computers tools and abilities more closely approaching those of human reasoning. Not only can these self-learning, self-modifying, self-writing programs learn with practice to get better at performing a particular task, they are able to change their fundamental structure to become faster and more effective at changing themselves to become faster and more effective. This has led to the current situation of computers becoming more intelligent at an increasing rate. Because these computers can rewrite their own programs, they are able to solve problems in ways that we sometimes cannot understand.

New feats of computer intelligence are proclaimed weekly, seemingly hourly. Computers can learn to play video games by reading the written instructions [5], read and evaluate legal documents [6], and scan online news reports and write accurate and informative newspaper articles [7]. They can look at images of dogs and determine their breed [8]. They can with increasing accuracy pick out faces in a crowd [9], and look at photographs of individuals and determine who those individuals are [10], what their emotional state is [11,12], and whether or not they are gay [13]. Computers can now diagnose skin lesions from photos as well as dermatologists [14] and recognize and transcribe speech better than teams of humans [15,16]. In radiology, computers are improving at a variety of image interpretation tasks, including determining bone age [17], finding fractures on plain films [18], and detecting interstitial infiltrates [19].

ARE THERE LIMITS TO COMPUTER INTELLIGENCE?

Given that computers are already in a real sense intelligent, will there be an ultimate limit to how intelligent computers will become? It is easy to find problems that will always be impossible for either humans or computers to solve, but are there problems that humans can solve that computers cannot?

The answer is unclear. There are two possible reasons why computers may not be able to fully duplicate human intelligence.

First, computers solve problems differently than the human brain solves problems. Computers function by a defined process called "computing"—this involves manipulating symbols, usually 0s and 1s. Although we do not know exactly how the human brain works, we do know that it does not work like that. It may be that there are types of problems that cannot be solved by the process of computation.

Second, computers are not now, and arguably will never be, conscious. It may be that some problems can only be solved by a conscious mind.

So it is possible that these two factors may impose limits on some aspects of computer intelligence. But then again, they may not. It may be that, even though computers do not "think" like humans, they may nevertheless be able to solve the same problems humans can using different methods—an airplane flies, but it does not fly like a bird flies. It is even conceivable that by using the process of computation, computers will be able to solve all types of problems better than conscious human minds can—and perhaps solve problems that human minds cannot.

AI AND THE FUTURE OF RADIOLOGY

Will computers eventually read all imaging studies as well as or better than human radiologists? When thinking about this, radiologists often point to particularly challenging problems—finding lobular cancer in dense breasts or comparing multiple postoperative spine MRIs with patient motion from different machines—and claim that computers will just never be able to figure these things out. That is an understandable reaction from a practicing radiologist, but it is like looking at a kindergartener and believing that, because she cannot add or subtract very well, she will obviously *never* be able to read an abdominal ultrasound. It assumes limits to computer intelligence that might not exist.

We are no longer talking about the kinds of systems that radiologists have worked with in the past—some early mammographic computer-aided detection that overcalled calcifications, undercalled masses, completely missed subtle asymmetries, and generally made a radiologist's life more difficult. We are talking about systems that are already on the horizon: computers that possess real intelligence, that can observe the world and learn from what they observe, and that steadily become more capable. If computers can do something now, they will only get better at it. If computers cannot do something now, they will probably learn how to.

What we will eventually see in radiology are diagnostic image interpretation systems that have read every textbook and journal article; know all of a patient's history, records, and laboratory reports; and have memorized millions of imaging studies. It may help to imagine these systems not as a collection of circuits in a console, but as an army of fellowship-trained radiologists with photographic memories, IQs of 500, and no need for food or sleep.

CONCLUSION

Life began on earth 4 billion years ago, evolved slowly, and remained comparatively small and simple until 500 million years ago when, in an event called the Cambrian explosion, the types and complexity of life rapidly increased to produce the variety that exists today. It is likely that we are at this moment in the midst of a similar explosion in computer intelligence. The advent of computers that can accurately interpret diagnostic imaging studies will upend the practice of radiology. The two currently unanswered questions are just how much upending there will be and how long it will take to happen. There are vastly differing opinions, from the apocalyptic claim that AI will make all radiologists extinct to the delusional assertion that computers will always merely assist—and never replace—radiologists. Both extremes are mistaken, but the truth is in the direction of the first.

Consider the fate of horses after the advent of machine power. In 1920 there were 25 million horses in the United States; in 1960 there were six million. Today there are 34,000 radiologists in the United States. Unless radiologists do things other than interpret imaging studies, there will be need for far fewer of them. This is a complex situation filled with unknowns, and events are moving fast. We need to figure how to deal with this coming change. And we need to do it in a hurry.

REFERENCES

1. Future of Life Institute. An open letter—research priorities for robust and beneficial artificial intelligence. Available at: https://futureoflife.org/ai-open-letter/. Accessed March 18, 2018.
2. Gibbs S. Elon Musk: artificial intelligence is our biggest existential threat. The Guardian. Available at:. https://www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat. Accessed March 18, 2018.
3. Vincent J. Watch Google's AI master the infamously difficult Atari game Montezuma's Revenge. The Verge. Available at:. https://www.theverge.com/2016/6/9/11893002/google-ai-deepmind-atari-montezumas-revenge. Accessed March 18, 2018.
4. Snow J. Computers learn to cooperate better than humans. Science. Available at:. http://www.sciencemag.org/news/2017/03/computers-learn-cooperate-better-humans. Accessed March 18, 2018.
5. Kaplan R, Sauer C, Sosa A. Beating Atari with natural language guided reinforcement learning. Stanford University. Available at: https://web.stanford.edu/class/cs224n/reports/2762090.pdf. Accessed March 18, 2018.
6. Lohr S. A.I. is doing legal work. But it won't replace lawyers. Yet. New York Times. Available at:. https://www.nytimes.com/2017/03/19/technology/lawyers-artificial-intelligence.html. Accessed March 18, 2018.
7. Keohane J. What news-writing bots mean for the future of journalism. Wired. Available at:. https://www.wired.com/2017/02/robots-wrote-this-story/. Accessed March 18, 2018.
8. Garun N. Amazon's image recognition AI can identify your dog down to its breed. The Verge. Available at:. https://www.theverge.com/2016/11/30/13799582/amazon-rekognition-machine-learning-image-processing. Accessed March 18, 2018.
9. Spice B. Finding faces in a crowd. Carnegie Mellon University. Available at:. https://www.cmu.edu/news/stories/archives/2017/march/faces-in-crowd.html. Accessed March 18, 2018.
10. Dormehl L. How machines learned to recognize our faces so well—and what's next. Fast Company. Available at:. https://www.fastcompany.com/3032386/how-machines-learned-to-recognize-our-faces-so-well-and-whats-next. Accessed March 18, 2018.
11. Andrade N. Computers are getting better than humans at facial recognition. The Atlantic. Available at:. https://www.theatlantic.com/technology/archive/2014/06/bad-news-computers-are-getting-better-than-we-are-at-facial-recognition/372377/. Accessed March 18, 2018.
12. Bartlett MS, Littlewort GC, Frank Mark G, et al. Automatic decoding of facial movements reveals deceptive pain expressions. Curr Biol 2014;24:738-43. Available at:. http://www.cell.com/current-biology/fulltext/S0960-9822(14)00147-X. Accessed March 18, 2018.

13. Wang Y, Kosinski M. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. J Pers Soc Psychol 2018;1149:246-57. Available at:. http://psycnet.apa.org/record/2018-03783-002. Accessed March 18, 2018.
14. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature 2017;542:115-8. Available at:. https://www.nature.com/articles/nature21056. Accessed March 18, 2018.
15. Emerging Technology from the arXiv. First computer to match humans in conversational speech recognition. MIT Technology Review. Available at:. https://www.technologyreview.com/s/602714/first-computer-to-match-humans-in-conversational-speech-recognition/. Accessed March 18, 2018.
16. Weinberger M. Microsoft's voice-recognition tech is now better than even teams of humans at transcribing conversations. Business Insider. Available at:. http://www.businessinsider.com/microsoft-research-beats-humans-at-speech-transcription-2017-8. Accessed April 18, 2018.
17. Lee H, Tajmir S, Jenny Lee, et al. Fully automated deep learning system for bone age assessment. J Digit Imaging 2017;30:427-41. Available at:. https://link.springer.com/article/10.1007/s10278-017-9955-8. Accessed March 18, 2018.
18. Dimililer K. IBFDS: Intelligent bone fracture detection system. Procedia Comput Sci 2017;120:260-7. Available at:. https://www.sciencedirect.com/science/article/pii/S1877050917324493. Accessed March 18, 2018.
19. Katsuragawa S, Ishida T, Ashizawa K, et al. Detection and diagnosis of interstitial lung disease. In: Qiang L, Nishikawa R, eds. Computer-Aided Detection and Diagnosis in Medical Imaging. New York, NY: CRC Press; 2015:169-85.

Robert Schier is from Radnet, Orinda, California.

The author has no conflicts of interest related to the material discussed in this article.

Robert Schier, MD: Radnet, 110 Ardith Dr, Orinda, CA 94563; e-mail: bobschier@gmail.com.