Rest assured, João, you are safe from artificial intelligence

Eliot Siegel, MD

Radiology will doubtless be one of the last medical specialties to be replaced by computers

Dr. Siegel is a Professor and Vice Chair, Research Information Systems, at the University of Maryland School of Medicine Department of Diagnostic Radiology and Nuclear Medicine, Chief of Imaging at the VA Maryland Healthcare System. He is also a member of the Editorial Advisory Board of Applied Radiology

My name is João Louro, and I am a first-year radiology resident from Oporto, Portugal. I decided to go into radiology based on my fondness for theoretical problems, informatics, and the need for a broad range of medical knowledge. Recently, I started learning about artificial intelligence (AI), and that has made me wary of the future of diagnostic radiologists. If everything is automated, will my work [consist of] just validating pre-made results and reports? Will I even be needed in the [age of] Big Data and AI?

I am just 25 years old with many, I hope, working years ahead of me. [Artificial intelligence has] even made me rethink my decision to go into this area. In the news I read the future of radiology is bright, but what about the future of radiologists? I am contacting you because I know you are a leading expert and are involved in these areas of research, and I want to know if you even recommend Diagnostic Radiology as a career at this point in time.

I have received dozens of emails, letters, and in-person queries in the past few years from residents, medical students, fellows, and radiologists in practice with the same concerns raised by João. I'm taking the opportunity to write this editorial to respond to each concern.

First of all, thank you, João, for permitting me to share your email. I can't blame you and your colleagues in radiology for being wary of your future or for your concerns about the potential of becoming merely a "validator of pre-made results" or of becoming completely replaced by machines. It's difficult to go a week without reading some new article about AI in medicine. Some of the world's authorities on machine learning, such as Professor Andrew Ng, of Stanford University, or Professor Geoffrey Hinton, a pioneer in artificial neural networks who divides his time between Google and the University of Toronto, have made some pretty scary pronouncements about radiology and radiologists.

Dr. Ng was quoted in *The Economist* as saying, "A highly trained and specialized radiologist may now be in greater danger of being replaced by a machine than his own executive assistant.[1] Dr. Hinton tweeted in October 2016, "We should stop training radiologists right now, in 5 years #deeplearning will have better performance."[2]

Meanwhile, Ezekiel Emanuel, credited as key architect of the Affordable Care Act, but with expertise in neither machine learning nor radiology, gushed in his recent address at the annual American College of Radiology meeting and in an article for the *JACR* that, "The most potent threat to radiology as a specialty is machine learning." In the same article, Emanuel wrote that, "It took 10 years for driverless cars to go from skepticism to reality. How long will it take for machines reading CT scans? Reading a CT or an MRI scan is much easier than driving. The image is not in constant motion at 30 or 65 mph with other cars, cyclists, and deer unexpectedly darting out."[3]

Continued on page 6

Continued from page 4

But let me be clear: Based on my 25-plus years of experience as a radiologist, as well as my years as a researcher and a proponent of machine learning in radiology, these dire and ominous predictions are absolute nonsense**,** and they are completely divorced from the reality of where we actually are with machine learning applications in radiology.

Let me explain why I say that. But first, let's define machine learning.

Machine learning is a subset of so-called "artificial intelligence," which is actually a general term used to describe the use of computers to perform tasks that previously required human intelligence, such as speech recognition, chess, and self-driving vehicles. Machine learning does not actually imply "learning" as humans generally think of learning or intelligence. Instead, it refers to a set of algorithms that are used to make predictions about "outcomes" when they are given new data based on patterns from previous data and previous outcomes. More data can be added over time, resulting in better predictions. Thus, a machine-learning algorithm can be said to "learn" to make even better predictions.

For example, using data such as nodule shape, size, density, location, age and smoking history from the National Lung Screening Trial (NLST), we can use regression statistics to create a simple formula to predict which future lung nodules are likely to be malignant. We can get an even better predictive model using more "advanced" hybrid computer algorithms and statistical methods using one of thousands of different "machine learning" techniques. The resulting formula or predictive model is typically so much more complex that it is saved as an algorithm rather than as a simple formula. Machine learning depends heavily on linear algebra; in fact, it can be thought of as a way to approximate the solution of a complex linear algebra problem. Fortunately for machine learning approaches, most biological and physics data patterns can be grossly simplified using this technique. Lin and Tegmark of MIT refer to this as "cheap learning."[4]

Second, one might ask why so much attention has been paid to machine learning in medicine and radiology when the technique has been around for more than 50 years. This is mostly due to a combination of the attention paid to the 2011 *Jeopardy!* game show match, which utilized IBM's Watson computer system, and the emergence of intelligent assistants such as Apple's Siri. However, the more immediate cause is related to the success of ImageNet, which combines advanced graphics cards and "convolutional neural networks" to push computer recognition of common objects such as dogs, cats, planes, and others in low resolution (typically 256x256 pixels) pictures to levels approaching or exceeding the speed and accuracy of humans. Interestingly, this was done by making each pixel a variable, with a total of about 66,000 variables. Given the speed of the processors, these pixels could be combined in intricate ways to "describe" increasingly complex structures and patterns.

You might conclude, as have some experts in human vision in machine learning, that this almost "magical" analysis could work with diagnostic medical images. In theory, it can. However, in actuality, such analysis is incredibly primitive at this point. Furthermore, it isn't clear whether we are one-tenth or one-billionth of the way there with regard to the ability of computers to recognize findings as well as radiologists. Importantly, the task of telling "what's wrong with this picture" in CT or MRI requires contextual knowledge that is far beyond the capability of today's most sophisticated computers. In fact, Koch and Tononi proposed this as a test for computer "consciousness" or, alternatively, as a test for general AI, rather than the narrow AI represented by a chess or speech recognition program. Additionally, no data exist to date to suggest that the success of ImageNet in any way extrapolates to interrelated series of images such as those of an MRI scan. So-called general AI is thought to be at least 20 to 100 or more years away, and it is difficult to imagine replacing a radiologist without this level of intelligence.

Another major challenge is the lack of annotated, massive datasets used to train the ImageNet machine learning algorithms. We have huge collections of images, but these are not annotated anatomically, functionally or pathologically. Even incredibly curated and purpose-built datasets, such as the NLST, have surprisingly minimal annotation of the pixel data, and the data that are annotated, such as the LIDC from the NCI, consist of about 1,000 cases, with only lung nodules annotated.

In the shorter term, daunting regulatory and medico-legal hurdles confront AI interpretation. To replace radiologists, an algorithm would presumably have to achieve the equivalent of passing the boards of all medical specialties and gaining acceptance by relevant imaging boards, including radiology and nuclear medicine. U.S. FDA clearance would either have to be based on testing thousands of capabilities of separate narrow AI algorithms, each requiring huge annotated datasets, or require convincing the agency that a program for general AI in diagnostic imaging had been created.

Even if a hypothetical "radiologist reading robot" from the distant future were teleported to 2017, and even if this robot could interpret images as well as today's radiologists, I think it could take as many as 20 years to get the datasets and testing required by the FDA for clearance---and that's not taking into account the FDA requirement to document the software development process. The "black box" nature of machine learning makes this requirement a deal breaker unless fundamental changes take place in the clearance process. Who would insure the company providing the software that would "replace" the radiologist? Who would be liable in a malpractice case? How would hospitals credential and privilege the software—would credentialing even be necessary? Radiology will doubtless be one of the last medical specialties to be replaced by computers; I don't see anything on the horizon to challenge that belief.

What hospitals can learn from IBM Watson Health's challenges

Ronald B. Schilling, PhD

Despite often being hailed as the next frontier of success for artificial intelligence (AI) technologies, the healthcare industry is actually proving strangely resistant to digital transformation. There's no better illustration of this than the collapse of the recent partnership of IBM Watson Health and The University of Texas' MD Anderson Cancer Center. What had started out as the promising debut of Watson's capabilities in recommending cancer treatment plans ended after an audit found that the project was poorly implemented and unfocused, and that the technology itself was running on dated information.

Although IBM claims it is proud of the work done through the project—in particular that the technology was found to be 90 percent accurate when supporting lung cancer treatment decisions—it is clear that much work remains before AI can play a meaningful role in a setting such as a hospital. But why is that? IBM's Watson has proven itself a formidable change agent in other industries, including oil and gas, hospitality and transportation. Why did it fall at the first hurdle in health care?

The answer is bigger than what can be learned from the pilot project at MD Anderson, which is just symptomatic of a broader problem. The answer lies in the cognitive technology itself and its most underestimated limitation – mimicking a physician's intuition.

Cognitive technology, as a rule, is supposed to be able to perform tasks that only humans, to date, have been able to do – only it's supposed to do it faster and more accurately. But when it comes to diagnoses and surgical planning, the mental math that IBM Watson so effectively executes in other industries is only half the battle. What it misses when it comes to managing something as intricate and nuanced as patient outcomes is the intuitive aspect of care.

In medicine, the breakthroughs, both big and small, that happen every day have historically been pioneered by doctors combining what they know with how they think. There are endless examples of surgeons using their instincts in the operating theatre to come up with creative solutions to medical problems. A surgeon in Tanzania recently chose to teach a non-doctor to do brain surgery because the latter had the "swagger and demeanor" of a surgeon. It worked. Pediatric surgeon Redmond Burke, MD, performed the world's first endoscopic vascular ring division, and later developed a series of minimally invasive tools to reduce risk in cardiac operations.

When looking at what technology is going to be able to complement this kind of environment, one must seek solutions that can bridge the cognitive and intuitive capabilities that define a good doctor – and garner the kind of clinical efficacy it takes to affect patient outcomes. Technology solutions designed to enhance a doctor's role with medical imaging, with diagnoses, etc., must tackle innovation as both an art and a science.

So, although IBM Watson may be able to win a game of *Jeopardy!*, it is still a long way away from creating the hospital of the future.

Dr. Schilling is President of RBS Consulting Group, Los Altos Hills, CA. He is also a member of the Editorial Advisory Board of Applied Radiology.

Finally, an even more daunting challenge exists. Despite widespread statements that radiologists serve only as a "commodity" that merely translates pixel data into reports, we actually do so much more than that. Most of us serve hundreds of functions at our outpatient centers or hospitals that we don't give ourselves credit for. We radiologists not only predict, but we also judge, explain, quality check, counsel, teach, discover, console, explore and create.

So, João, my strong advice is to remain in your diagnostic radiology residency and to embrace a whole new emerging array of machine learning-based systems that will make you more effective, efficient, and safe in your profession. These algorithms will help us all with the mundane, time-intensive, and error-prone tasks and free us up to use our judgment, common sense, and superior intellect in ways that we are not able to today.

Based on what I know today, if I were still in medical school I would not hesitate to choose radiology as the specialty with the most exciting opportunities for collaboration with computers to discover and pave new frontiers in research and clinical care.

Indeed, until general AI arrives to supplant the human tasks that require real human intelligence, I assure you, João, radiology and radiologists are safe.

References

1. Morgenstern, M. Automation and anxiety. *The Economist.* https://1.www.economist.com/news/special-report/21700758-will-smarter-machines-cause-mass-unemployment-automation-and-anxiety. Accessed August 7, 2017.
2.Mukherjee, S.. A.I Versus M.D. *The New Yorker.* April 3, 2017. http://www.newyorker.com/magazine/2017/04/03/ai-versus-md. Accessed August 7, 2017.
3.Chockley, K. A. The end of radiology? Three threats to the future practice of radiology. *Journal of the American College of Radiology.* 2016; *13*(12):1415-1420.
4. Lin, H. W. Why does deep and cheap learning work so well? *Journal of Statistical Physics.* 2016; 1-25. Published online July 21, 2017. https://arxiv.org/pdf/1608.08225.pdf. Accessed August 8, 2017.