Is AI in Healthcare Doomed, or Destined for Greatness?

by Bijan Elahi
Eindhoven, The Netherlands

Healthcare costs are on the rise, worldwide. The majority of the people on this planet do not have access to quality healthcare. Medical errors are the third leading cause of death in America [Ref. 1]. Artificial intelligence (AI) can help.

Fully automatic and autonomous medical systems are already released and in use. Nurses and doctors have adopted this technology to reduce manual work, and to provide more accurate service and impactful interventions to patients. A survey of 142 information technology (IT) professionals in healthcare in 2018 showed that 77% are using AI to support clinical decisions, while 55% use AI/machine learning (ML) to make earlier diagnoses of diseases [Ref. 2]. According to Forbes magazine, virtual round-the-clock nurses could save $20 billion annually [Ref. 3]. AI help in administrative tasks, such as transcriptions, could save $18 billion annually. AI in robotic-assisted surgery could reduce complications by 21% and reduce hospital stays, saving $40 billion annually. An analysis by Accenture says that key AI applications in healthcare can create $150 billion in annual savings for the U.S. healthcare economy by 2026 [Ref. 4]. To put the magnitude of this in perspective, $150 billion is enough money to train 600,000 physicians[1], or build 670[2] fully equipped hospitals.

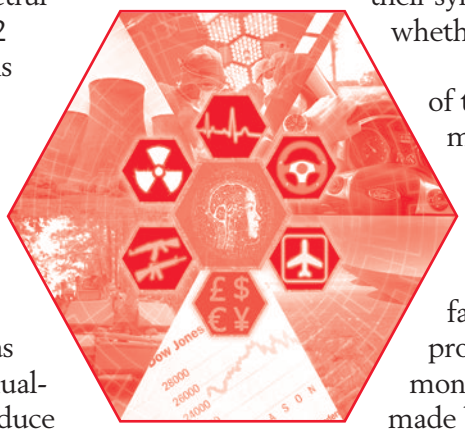

Performance

AI has been shown to outperform humans in numerous studies. In 2018 in Beijing, China, an AI software named BioMind was pitted against 15 top doctors in China in diagnosing brain tumors and predicting the expansion of brain hematomas. In two rounds of competition, BioMind outperformed the physicians every time [Ref. 5].

At one emergency call center, the operators triage callers to determine the most urgent suspected myocardial infarctions (MI) by listening to the caller's tone of voice, breathing, verbal and non-verbal patterns of communication. At this center, AI was able to identify MI victims with a 93% accuracy, vs. 73% accuracy rate for human operators. [Ref. 6]

AI has had one of the most significant impacts in radiological image processing. According to *Forbes* magazine, AI can analyze 3D radiological images up to 1,000 times faster than what is possible today with human analysis. [Ref. 3]

Access

In the U.K., the National Health Service (NHS) uses an automated smartphone triage system called "Babylon GP at Hand," through which people can get advice about their symptoms as well as a recommendation of whether they need to go to a hospital.

According to *Forbes* magazine, 15.5% of the global population is afflicted with mental illness. More than half of cases of mental illness go untreated, and every 40 seconds one person dies from suicide [Ref. 7]. AI can diagnose mental illness by analyzing linguistic patterns, facial expressions and tone of voice and produce a diagnosis of mental illness three months before a formal diagnosis is typically made by a trained mental health practitioner.

Currently, an AI product called "IDx-DR," a fully autonomous system, evaluates retinal images of diabetes patients for retinopathy and decides whether the patients should be seen by an ophthalmologist. This is the world's first fully autonomous FDA-approved AI system.

What is AI?

Currently, there is no consensus on what "AI" is. In general, it is believed that AI is a simulation of human intelligence. Certain aspects of human intelligence such as facial and speech recognition, reasoning and learning are considered to be in the domain of AI.

Many people use AI today whether they are aware of it or not. Chatbots that interact with you on commercial websites and automated telephone answering systems that recognize your speech and process your calls are all AI applications.

Grades of AI

When we speak of AI today, we refer to Artificial Narrow Intelligence (ANI), also known as Applied AI. These are narrow areas where computers can perform on par with humans, or sometimes better than humans. Common examples include Google Translate, Siri/Alexa, image recognition and medical diagnostics.

[1] Based on the median cost of medical school in the U.S. in 2013, $250,000.
[2] Based on the cost of an average-sized, 150-bed hospital in the U.S., $224 million

Figure 1 — Example of Captcha.

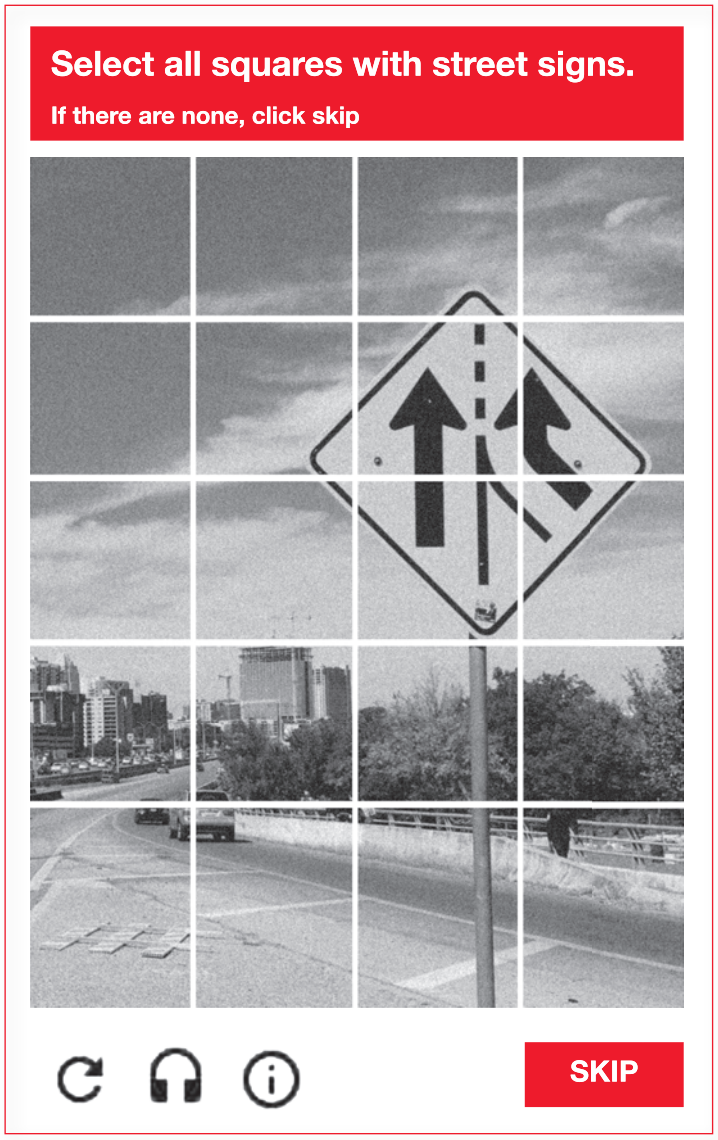

Figure 2 — Example of a Graphic Test for Human Authentication.

There is a consensus that at some point in the future computers will be able to perform as well as humans in all areas of human endeavor. In these cases, it would not be possible to distinguish AI performance from human performance. This is called Artificial General Intelligence (AGI). There are science fiction examples of AGI. The movies *Ex Machina* (2015) and *Morgan* (2016) are two such examples.

Since AI outperforms humans in *some* areas even today, when AGI is achieved, we can conclude that *overall* AI would cognitively outperform humans. This is called Artificial Super Intelligence (ASI). However, ASI is commonly understood to mean superior cognitive performance in *most* areas.

Figure 3 — Empathy

AI Strengths and Weaknesses

AI is particularly suited to data-rich domains. Advances in technologies such as wearables, genetic tests and new biomarkers are creating an explosion in the amount of information being presented to medical professionals, with the expectation that the medical professionals will integrate this information into accurate diagnostics. This is not humanly possible. AI, on the other hand, can process huge quantities of data and assist medical professionals.

AI is weak in associative thinking and in integrating widely diverse information. AI is also weak in working with seemingly contradictory, wrong or missing information. You may be familiar with "Captcha," which are tests that some websites apply to ensure you are a human, not a machine. Figure 1 is an example of a distorted set of characters that would be difficult for AI to decipher, but relatively easy for humans. Figure 2 is an example of a graphic challenge where a computer may not be able to distinguish which squares contain street signs, but a human can easily do so.

A human doctor can listen to a patient describe their symptoms incorrectly — giving inconsistent, incomplete and contradictory information — and still make a correct diagnosis. A human doctor can do this by dismissing clearly wrong input, filling in the missing information and creating consistency by correlating various information such as the smell of the patient's breath and the posture and/or look of the patient.

Another area where AI is weaker than humans is empathy. We humans have evolved to be able to mirror the feelings of another human. It would be difficult for AI to tell whether the person in Figure 3 is angry or concentrating. But a human in her presence would more readily be able to feel whether she is angry or just concentrating.

Fully Human Decision Making	Machine-assisted Human Decision Making	Fully Autonomous Machine Decision

Figure 4 — Spectrum of Human-to-AI Decision Making Control.

AI in Healthcare

In most areas of AI application, performance is the main goal, but in healthcare, patient safety is paramount. Therefore, spectacular progress is not as common in healthcare applications of AI. Some of the areas where spectacular performance is seen are:

- **Autonomous vehicles** — AI can not only decrease traffic deaths but operate vehicles more efficiently to save fuel and lower emissions
- **Fraud detection by credit card companies** — AI can assist in detecting fraud attempts that are becoming increasingly sophisticated
- **Crime fighting** — AI can predict areas and times where there is high likelihood of crime; also, in the justice system AI can predict recidivism and improve outcomes by releasing people who it predicts are less likely to commit crimes in the future
- **Aviation** — AI can fly damaged aircraft in the best possible way to land safely
- **Finance** — High-speed trading has been around for years via algorithms; AI can improve these outcomes via machine learning

It should be noted that in the healthcare field, *performance* and *outcomes* are not the same thing. For example, AI may outperform human physicians in certain narrow areas, but the ultimate outcome to the patient's health may be not as good as when a human physician treats the patient because a human physician can incorporate a lot of peripheral information that may not be available to AI. A physician may ask a patient about their emotional state; their happiness in their marriage or their job; notice an awkward gait, or bad breath. AI may not have access to all of this peripheral information.

Safety Risks of Using AI

AI's influence in technology falls on a spectrum. Figure 4 illustrates this range, from fully human-controlled decision making to fully autonomous machine decision making.

Risk is present all along this spectrum. Humans make mistakes too, but humans are more comfortable with the risks of status quo. The old adage "better the devil you know than the devil you don't know" highlights this mentality. Assuming the risks of fully human decision making is the *baseline*, also known as *the state of the art*, according to ISO 14971, let's examine the differential safety risks of introducing AI.

Case 1 – Machine-assisted decision making. In this case, humans are the final arbiter; the machine makes suggestions or provides information to help the human. The most common example of this is GPS navigation. Most people have used GPS navigation at some point. Because GPS navigation is correct and reliable in the vast majority of cases, humans gradually relinquish judgment and trust GPS implicitly. There are two risks associated with this:

1. GPS (or the machine assistant) is not always correct. You may have seen reports like the one found in Figure 5.

What happened here was an implicit and unexpected transfer of risk from the human to the machine. That is, although the manufacturer and society have presumed the human is in control and the final arbiter, the human has implicitly transferred control to the machine.

In medical technology, this can be seen in machine-assisted diagnostics. For example, in capsule endoscopy, a patient swallows a medical device in the form of a capsule that contains cameras and lighting. The capsule takes pictures of the digestive tract as it travels from the mouth through the anus. This trip takes many hours and results in many thousands of photos. The advantage is that parts of the body that, to date, have not been accessible can now be examined. The problems is that to review all the images would take hours — perhaps even days — of a doctor's time.

To help with this, software has been developed to examine the pictures and identify those images that are suspect. This alleviates a significant burden on the physi-

Woman Follows GPS Into Lake

Figure 5 — Implicit Transfer of Risk to Machines.

cian. In the future, AI could be used to identify the cancerous polyps and bring them to the doctor's attention. Let's imagine that AI is so good that the doctor develops enough trust that he or she stops verifying the AI diagnostic. In this example, the risk is implicitly transferred to the machine, which was not the intent of the manufacturer.

2. Human skill is lost. Reliance on AI, or automation in general, can lead to complacency and a loss of human skill. Going back to the example of the GPS navigation, people who use GPS frequently, often find that their skills in navigation using the traditional methods, such as paper maps and landmarks, gradually degrades. Airline pilots have access to automatic landing systems, but to maintain their skills they often do manual landing of the aircraft. Not using AI when it presents tremendous benefit, convenience or comfort takes a great degree of discipline, particularly because it is human nature to seek the easiest solutions. Additionally, external pressures can be a factor. Consider a hospital that is trying to maximize profits. They would pressure the doctors to produce more results with fewer resources. You can imagine that under such pressures, a physician might tend to trust the AI diagnostics and rapidly approve capsule endoscopy results.

There may be even more risk due to unanticipated outcomes. What makes AI special is its ability to learn, adapt and improve itself. The downside of this is that AI can take unexpected turns that can be incomprehensible. *Forbes* magazine reported in July 2017 that Facebook AI chatbots had learned to communicate with each other and, in the process, had created a new language that was not comprehensible to humans [Ref. 8]. We need to be prepared for more such events to happen.

Can We be Safe with AI?

In medical technology, safety is defined as freedom from unacceptable risk. It is widely accepted that 100% safety is unachievable. Risk management strives to assess the level of risk and balance it against the benefits of the medical technology. Medical systems offer certain benefits, such as alleviating pain, restoring function and extending life. But they also present risks. The question is whether the risks of medical systems that deploy automation and autonomy are outweighed by the benefits provided.

The risks of traditional medical systems, such as X-ray machines, are well understood. But AI systems are new and techniques for estimating of their risk need to be developed. Regulatory bodies also need to formulate methods to evaluate AI medical systems before approving them for commercialization.

An interesting human phenomenon is that people are less forgiving of AI than they are of humans. We expect infallibility of AI, but we tolerate some degree of error from humans. Because of this, people might abandon AI in healthcare technology much sooner if AI makes a mistake. In complex domains there will be mistakes. Surprisingly, this aversion to AI could ultimately mean worse outcomes would be accepted by the community by favoring humans over AI.

Another psychological phenomenon at play is resistance to loss of control. Physicians would likely not be comfortable with giving up control of their patient's treatment to a machine — even if the machine shows superior performance. This is nothing new. Past history has shown that every time automation threatened to take away some control, it has met with resistance.

Early in the 20th century, when elevators were first introduced in high-rise buildings in New York City, every elevator had an operator to drive the elevator because people were afraid of entering the elevator unless a human was driving it. In 1945 the New York City elevator operators went on strike. Business in the city was impacted significantly. For a week, no one dared ride an elevator. To alleviate the fear, every elevator was equipped with a red telephone that people could pick up and talk to someone who would guide them to safety. The elevator companies also added a big red "STOP" button to elevator control panels. That way, passengers could push the big red button and stop the elevator. Of course, this did not reduce the likelihood of an elevator crash, but it did make people feel that they had some degree of control, and that made them feel better. Gradually, people got used to driving the elevator themselves and their fears subsided.

Other Challenges with the Use of AI

There could come a time when AI is in charge of making decisions such as when to remove life support from a terminal patient. Or, AI may decide to keep a patient on artificial life support, prolonging their agony. These situations would raise ethical issues. Currently, medical technology is not prepared to handle the ethical issues of allowing AI to make such decisions.

Another potential area of concern is legal liability. If an AI-based medical device causes the death of a patient, who is liable? We learned that unsupervised machine learning is a situation where a human has not programmed the machine to do something — there is no algorithm. AI has learned by observation. How would a court of law hold AI accountable for its mistakes?

Wisdom

According to a study by Frost and Sullivan [Ref. 9], the momentum for AI usage in the healthcare industry is building up and is likely to exceed an annual growth rate of 40% by 2021. It's clear that AI is just beginning and there's no stopping it.

Do not be afraid. AI is just another new shiny tool. Every time a new tool proliferates, fears rise. For example, we use Excel spreadsheets to store and analyze data. Anyone who uses Excel knows that it can produce wrong answers. This could be due to no fault of Excel, but due to the erroneous use of the tool. Should we not use Excel? Mobile phones offer a great convenience. But sometimes calls are dropped. Should we not use mobile phones?

The wisest thing to do is to understand the weaknesses of AI, understand its potential for mistakes, and put it into context. If an autonomous car gets into an accident and someone dies, compare the rate of fatal crashes of autonomous cars with the rate of fatal crashes of human-operated automobiles.

The best approach would be to combine human and AI in an optimal way. The international Symposium on Biomedical Imaging did an experiment to compare the accuracy of detection of metastatic breast cancer. The machine's accuracy was 92.5%. The humans' accuracy was 96.6%. But when they combined human and machine, where the humans were given the machine's diagnosis, the combined accuracy went up to 99.5% [Ref. 10].

Table 1 – List of Terms and Acronyms

Term	Definition
AI	Artificial Intelligence
AGI	Artificial General Intelligence
ANI	Artificial Narrow Intelligence
ASI	Artificial Super Intelligence
FDA	Food and Drug Administration – U.S. Regulatory body
GPS	Global Positioning System
ISO 14971	International Standard — Application of risk management to medical devices
IT	Information Technology
MI	Myocardial Infarction
ML	Machine Learning
NHS	National Health Service (U.K.)
GPS	Global Positioning System

About the Author

Bijan Elahi is an author and a global product risk-management consultant for Medtronic. He is also an educator and lecturer at the Eindhoven University of Technology and the Delft University of Technology.

References

1. Makary, Martin A. "Medical error – third leading cause of death in the US," *BMJ* 2016;353:i2139, 2016.
2. Sullivan, Tom. "3 charts show where artificial intelligence is making an impact in healthcare right now," *HealthcareIT News*, December 21, 2018. Retrieved from https://www.healthcareitnews.com/news/3-charts-show-where-artificial-intelligence-making-impact-healthcare-right-now
3. Marr, Bernard. "How Is AI Used in Healthcare — 5 Powerful Real-World Examples That Show the Latest Advances," *Forbes*, July 27, 2018. Retrieved from https://www.forbes.com/sites/bernardmarr/2018/07/27/how-is-ai-used-in-healthcare-5-powerful-real-world-examples- that-show-the-latest-advances/#6fa897dd5dfb
4. Accenture. "Artificial Intelligence: Healthcare's New Nervous System," 2017. Retrieved from https://www.accenture.com/_acnmedia/PDF-49/Accenture-Health-Artificial-Intelligence.pdf
5. Wang, Xiaodong. "AI defeats elite doctors in diagnosis competition," *The Telegraph*, July 24, 2018. Retrieved from https://www.telegraph.co.uk/china-watch/technology/artificial-intelligence-in-medicine/
6. Vincent, James. "AI that detects cardiac arrests during emergency calls will be tested across Europe this summer," *The Verge*, April 2018. Retrieved from https://www.theverge.com/2018/4/25/17278994/ai-cardiac-arrest-corti-emergency-call-response
7. Marr, Bernard. "The Incredible Ways Artificial Intelligence is Now Used in Mental Health," *Forbes*, May 3, 2019. Retrieved from https://www.forbes.com/sites/bernardmarr/2019/05/03/the-incredible-ways-artificial-intelligence-is- now-used-in-mental-health/#47ed6d05d02e
8. Bradley, Tony. "Facebook AI Creates Its Own Language In Creepy Preview Of Our Potential Future," *Forbes*, July 31, 2017. Retrieved from https://www.forbes.com/sites/tonybradley/2017/07/31/facebook-ai-creates-its-own-language- in-creepy-preview-of-our-potential-future/#5e80d905292c
9. Frost & Sullivan. "From $600 M to $6 Billion, Artificial Intelligence Systems Poised for Dramatic Market Expansion in Healthcare," January 5, 2016. Retrieved from https://ww2.frost.com/news/press-releases/600-m-6-billion-artificial-intelligence-systems-poised-dramatic-market- expansion-healthcare/
10. Wang, Dayong, et al. "Deep Learning for Identifying Metastatic Breast Cancer," Cornell University, June 18, 2016. Retrieved from https://arxiv.org/abs/1606.05718