Diego Carrasco

Dr. Hoover

Comp 211

25 November 2019

Artificial Intelligence: An Innovative Technology for a Vital Industry

Everyone gets sick eventually, and the vast majority of sick people inevitably end up in some kind of medical facility: hospitals, clinics, or emergency rooms. And although these places are generally assumed to be safe, this is statistically not always the case. According to an eight-year study by John Hopkins patient-safety experts, more than 250,000 deaths per year can be attributed to medical errors, putting medical errors as the third leading cause of death in the United States after heart disease and cancer (McMains). Despite their prevalence, medical errors are not actually recognized by the Center for Disease Control (CDC); instead, deaths due to medical error are classified by the specific cause of death. Medical errors include deaths caused by medication errors, environmental infections from medical centers, poor patient follow-up plans, and delayed or missed diagnoses (Carver). However, advancements in technology could help reduce the problems. Artificial intelligence has been in development since the 1950s and is currently one of the fastest-developing technologies in the world. Although movies like *I, Robot* and *Terminator* have given a negative portrayal of the technology, it has largely been used in various other industries to increase both productivity and safety. The automobile industry, for example, has been developing programs that use AI to keep cars in their lanes and prevent or minimize the damage from accidents (Schroer). Using this and applications from other industries, artificial intelligence could reduce human error, and subsequently medical errors, in the healthcare industry.

Although artificial intelligence seems like a relatively new concept, the idea has been around since the seventeenth century. In the mid 1600s, philosopher and mathematician Rene Descartes theorized about two kinds of machines: one that could learn to perform a specific task and another that could adapt to perform a variety of tasks, much like humans (Sraders). Although Descartes was restricted by the technology of his time, he had actually developed two core ideas of artificial intelligence: specialized and generalized AI (Sraders). After the invention of the first programmable computer in 1938, computer scientists began thinking about what functions they could program into computers (When). Over the next decade, computer scientists worked on creating calculating programs with these new computers, but one scientist saw potential far greater than calculating numbers. In 1950, after theorizing that computers could be programmed to think in his now-famous paper "Computing Machinery and Intelligence," computer scientist Alan Turing created a test evaluating the intelligence of a machine (Wilson 39). Known as the Turing Test, this assessment states that if a machine could give responses similar enough to a person's as to be undifferentiable from an outsider's perspective, then the machine can be considered intelligent. The test was never meant to be passed; rather, it was created to theorize about the kind of programming required for a machine to be considered intelligent (Wilson 44). The technology would progress through the remainder of the century and be organized into schools of thought, further complicating an already complex area of science.

Although "intelligence" may seem like a straightforward descriptor, the blanket term does not do justice to the complexities of modern artificial intelligence systems. Most AI systems are classified as either strong or weak AI (Sraders). Weak AI systems are supervised programs designed to simulate human thought and interaction without it actually comprehending the interaction. Siri and Alexa, for example, may appear to understand commands and requests, but

they are really just programmed with responses to certain key words and phrases, and this simulates intelligence. On the other hand, strong AI systems are almost completely unsupervised and usually involve processing data (Sraders). These systems are best known for teaching themselves new processes; in 2013, for example, a strong was programmed to AI learn to play seven Atari 2600 games and ended up breaking several game records and even besting a human expert in three of the games (Mnih 1). AI systems are also classified by function. Specialized AI systems are generally weak AI systems programmed to perform one specific task – news feeds, for example – and are the most common type of AI (Sraders). Generalized AI systems are strong AI that learn to perform multiple tasks; these are primarily the AI systems that have made headlines as major technological breakthroughs in recent years. All these kinds of technology – especially weak, specialized artificial intelligence systems – could be used to mitigate and reduce medical errors in the healthcare industry.

Although the United States has some of the most advanced medical technology in the world, currently, AI technology is not yet advanced or integrated enough to reduce a significant number of medical error deaths (McMains). The medical errors in question are preventable outcomes of medical care that end with the death of a patient or patients, which differs from errors that do not result in death (Carver). This broad category covers a wide variety of mistakes such as medicinal overdoses, misdiagnoses, undertreatment, overtreatment, surgical injuries, infections, and other errors. Medical errors can occur anywhere, but the fatal ones are most likely to happen in operating rooms, intensive care units, and emergency rooms than anywhere else due to the highly stressful environments those areas tend to have (Carver). However, medical errors are often not caused by individual recklessness. They are most often caused by environmental or systemic errors which cause human mistakes or failure to prevent human error. For example,

most hospitals keep full-strength drugs stocked, and these are fatal to patients unless properly diluted to specific amounts that vary with each patient's age, weight, sex, family history, and current medical status (Carver). Unfortunately, it is very easy to let an undiluted or under-diluted sample slip into a patient's IV bag, and this could prove fatal if it is not caught in time. Despite these oversights, instances of human error have certainly led to severe medical error. "Anchoring bias" is the propensity or tendency to depend far too heavily on the first piece of evidence when making decisions, and this bias is a major source of misdiagnoses in the medical field (The Anchoring Effect, Carver). However, when such errors occur, blaming the individual does not solve the problem for when the next physician makes the same error. In addition to taking measures to prevent medical errors, medical practices should seek to make it more difficult to make errors, and this could be accomplished through artificial intelligence.

There are various uses of artificial intelligence technology in other industries that could be applied to the healthcare industry to prevent medical errors, and some of the most common uses of the technology occur in businesses. Most people are familiar with the 24/7 customer service AI commonly found in businesses, but AI is far more involved than this. According to a Harvard Business study, the most common use of AI in business was in robotic process automation, also called RPA (Davenport). This weak, specialized AI has a variety of functions including recording and transferring data from emails and phone calls, replacing credit cards, and reading legal documents to extract pertinent data. Because it is the least expensive and simplest AI, RPA generally has the quickest and highest returns for a business, and this type of AI could be easily integrated into healthcare to organize records or help with payment issues much more efficiently than by hand. Another common function of AI in business is cognitive insight, which detects patterns in large quantities of data and forms conclusions from them (Davenport). Among

other things, this AI is generally used in identifying customer patterns and spotting identity theft. It also is able to improve over time by retaining previous information, unlike RPA. Cognitive insight AI could be useful in conducting studies about specific diseases or illnesses and would be much more efficient and effective than a human going through the data. Cognitive engagement is the least used AI in businesses, but it is also the most advanced. This strong, generalized AI can engage customers in conversation using natural language, help employees choose healthcare plans using the individual's health and past medical care, IT support, and various other functions requiring research, data searching, and forming more complex solutions than cognitive insight AI (Davenport). Although it is the most advanced technology of the mentioned selection, it sometimes cannot keep up with customer demands, forcing some companies like Facebook to restrict their cognitive engagement AI chatbots to specific topics.

The automobile industry is also filled with technologically-innovative AI systems that could be used in the healthcare industry. Self-driving cars are exemplified quite often when discussing artificial intelligence, but the basic technology behind the specialized machines is astounding. In Boston, a company called nuTonomy has developed nuCore, which allows vehicles to drive through and learn from complex traffic scenarios; recently, they have teamed-up with Lyft to test drive their automobiles with Lyft customers in Boston's Seaport District to much success (Schroer). At its core, this technology is a highly advanced learning program designed to make life-changing decisions rapidly by assessing the most efficient path to save the most human lives, and this could be practical in a stressful decision-making environment like an emergency room or trauma center. AutoX is using a similar technology to automate grocery delivery with trials currently in San Jose, California and plans to expand outward from there. This idea could be expanded to medication delivery or even first aid transport for older patients

or those with disabilities. Waymo, a former technological substituent of Google, has developed cameras with a 360-degree view that can detect pedestrians, cyclists, vehicles, road work, and other obstacles up to 300 yards away, which shows how far detection technology has come (Schroer). This technology could be used by a doctor or nurse to analyze patients and look for specific external symptoms to suggest or eliminate conditions. On the manufacturing side, DataRPM has developed a technology that analyzes a machine's patterns of behavior and predicts future breakdowns or machine failures, which saves both time and money. Additionally, Rethink Robots have created robot AI designed to collaborate with humans in car manufacturing, combining human supervision and imagination with the efficiency of robotics (Breunig). The robots are programmed by being physically walked-through the process, which saves money on programming costs and increases efficiency of tasks which are not fully automatable by 20% (Breunig). Overall, the automotive industry provides some of the most creative advancements in artificial intelligence technology which could improve the healthcare industry.

Although healthcare is primarily focused on the medical professionals and their expertise, artificial intelligence is already at play behind the scenes. These systems can perform a variety of tasks but are primarily used for accessing specific material from a sea of medical knowledge like treatments, medical studies, related medical cases, and survival rates (Kennedy). Much like the aforementioned AI used for cognitive insight, this AI processes large quantities of information for whatever it has been programmed to see – trends, symptoms, diseases, things of that nature – and draws conclusions from the data it gathers. It also retains and stores the information for future reference, thus allowing it to "learn" in the most basic sense (Kennedy). For example, a 2006 study of epidemiology – the study of incidence, distribution, and control of diseases – tested the boundaries of artificial intelligence algorithms against previous methods (Flouris 167).

In the past, epidemiology was resigned to purely population-level analysis because the analyses were all done by individuals or teams of people. After testing out a new AI algorithm, the researchers found that the AI was not only able to sift through all the old information quicker than normal, but also was able to expand and apply itself to even larger populations, thus evolving epidemiology to unprecedented micro and macroscale capabilities. This innovation allowed further studies on diseases and illnesses that progressed medicine and medical treatments beyond what was previously capable.

Given the current success of artificial intelligence in the healthcare industry, additional implementation of the technology would further reduce medical errors, and this process can begin as soon as a patient enters into a medical facility. According to a survey by the Medical Group Management Association, patients spend twenty minutes on average in the waiting rooms of medical clinics (Heath). Instead of letting this time go to waste, it would be beneficial to include an AI system on a tablet or computer similar to Siri, Alexa or perhaps even a customer service line that would record and process patient symptoms while in the waiting room. The attending doctor or nurse in turn can have a general idea of the symptoms before the patient arrives, allowing the patient to get faster service than before. Beyond just recording the information, the AI could also cross-reference the results with matching symptoms and hypothesize a condition before the doctor arrives. Of course, this could easily turn into anchoring bias, but it would be foolish for a doctor to solely trust the information of a program above his or her years of medical training. This technology, although it would not replace a doctor's expertise, may reduce diagnostic errors by narrowing down the symptoms to a single or handful of conditions before beginning any tests, which would eliminate or at least reduce diagnosis-

related medical errors like failure to diagnose, delayed diagnoses, frivolous tests and incorrect diagnoses (Carver).

Another potential use of artificial intelligence is for aforementioned medicinal drug overdoses. Instead of leaving this process to human error, it would be better to automate drug dilution with a weak AI system programmed to dilute drugs to non-deadly portions based on the age, weight, symptoms, sex, and medical history of the patient, all of which could be obtained from the waiting room survey previously suggested. The human brain is incredibly powerful and dwarfs all modern computers in processing power and cognition. However, computer programs are now able to perform calculations and store information far better than human brains when performing specialized tasks – chess-playing, for example. In 1997, former chess world-champion Gary Kasparaov was beaten by IBM's "Deep Blue" supercomputer, which was a massive upset since Kasparov had scoffed at the idea a decade earlier (Gibbs). Since then, AI programs have consistently bested the highest-ranking chess players and are even beating other highly advanced AI chess programs. This demonstrates that among one of the most contemplative and analytical activities, specialized AI systems are better than human brains, and this idea could be applied and used to automate tasks like drug dilution. This eliminates the human aspect of the process and almost completely eliminates the probability of human error when diluting medical drugs. Although this is a specific situation, the general attitude toward implementation of AI in the healthcare industry should be making it more difficult to commit error instead of reducing error, and automating some easily mistaken tasks would accomplish this goal. This outlook combined with the safety net of AI should make medical facilities much safer.

DataRPM has developed a program that analyzes machine behavior and looks for mechanical failures before they can happen for the automotive industry (Schroer). Although the technology may not be fully available yet, it is theoretically possible to apply this technology to surgical rooms to look out for improper incisions, monitor patient health, and maintain a general watch over the surgery while the doctors and nurses focus on the more minute details. Furthermore, the technology developed by Waymo that identifies foreign objects in the road could be readjusted to search for surgical errors which may escape a human's perception. Combined with the DataRPM technology, the resulting program could save hundreds of lives from surgical injuries, wrong-site surgery, and maybe even help with identifying medical problems. Again, this would not replace doctors and their expertise, but would serve as a tool to make their work more efficient.

Despite the advantages of using artificial intelligence in not only healthcare but all industries, there are still quite a few obstacles holding AI back. One of the primary arguments against AI is that it is difficult to trust such an intrusive technology without fully understanding its decision-making processes – including potential biases – or the intentions of the developers, especially with something as sensitive and valuable as healthcare (Rossi 128). There is some merit to this argument because users have a right to know what they are using and who specifically designed it, and the best way to counter it is through transparency. Users want AI systems that can be explained thoroughly and appear to have ethical and moral stability that keep the AI free of bias and morally good (Rossi 128). To answer the call, IBM, one of the leading companies of AI technology, has published a set of transparency and trust capabilities for their AI system OpenScale designed around "explainability, fairness, and traceability" (Rossi 130). This publication explains the functions of the AI and its capacity, denotes the measures taken to

eliminate bias from the system, and provides a line to the developers and researchers to report any apparent biases with a goal of tweaking the program to prevent such occurrences. IBM's initiative not only encourages trust between the company and the consumers, but also helps refine the technology using the users' inputs, further strengthening that relationship of trust. Additionally, AI producers are also beginning to credit the developers of the technology with names and faces so that they no longer exist as anonymous creators of an intrusive technology. These and other measures focused on trust should help mitigate the worries people have about trusting the technology.

Another argument against AI technology is one of ethics: should an AI system malfunction, who would be bear the blame? Would it be the owner of the machine? The last person to use it? The developer of the AI? This argument is typically applied to decisions self-driving cars may have to make on the road, but it could be applied to any AI technology where the system has some degree of independence from humans. And although the answer is not as clear for autonomous automobiles, it is a bit less murky when applied to healthcare. Where self-driving cars are naturally placed in situations where they are forced to make split-second analytical decisions regarding peoples' lives, healthcare AI should never have the full autonomy to make those kinds of decisions; it should always be accompanied by a doctor or nurse to make sure the AI maintains a certain level of subservience to the medical professionals. This keeps the responsibility on the medical professionals and reduces the AI to an accessory role and not an executive one. However, this raises yet another question: how will older, less technologically-oriented physicians adapt to the technology? This dilemma opens the opportunity for a unique teaching relationship between older and younger doctors. Doctors learn from each other and pass their knowledge and experience down to the next generation of doctors. However, pairing older

physicians with younger pupils proficient with newer technology would allow the student to become the teacher and help ease the transition into a more technologically focused industry. Although not every experienced doctor will accept the technology and will argue that they do not need the help or precautions, it is still possible to integrate the technology using the younger generation of doctors. If medical schools begin equipping students in medical AI technology, a new generation of doctors will enter the medical field with knowledge and experience on how to use it. Eventually, medical professionals will all be capable of using artificial intelligence in their practice, and the medical industry will continue to progress.

One of the biggest counterarguments to artificial intelligent is the fear that AI will eventually become self-aware and hostile toward humanity. This is the plot of several movies including *Terminator 2: Judgement Day*, *Avengers: Age of Ultron*, *I, Robot*, and *The Matrix*. In general, these movies' plots revolve around an AI that evolved beyond what its human creators imagined, saw humanity as an enemy of some sort and sought to destroy or enslave humanity. Although this idea makes for some fantastic movies, it is impossible by its own nature. All AI technology that has been produced has reflected human thinking; essentially, AI think like humans because that is all humans know to program – this is reflected in the Turing Test where intelligence in a machine is measured by its ability to respond to a human in a human-like way (Dobelli). AI is always used to accomplish human objectives just like a human, but with fewer errors and faster processing speed. All AI reflect this human way of thinking and will eventually reflect human goals, aspirations, and morality because that is all humans know how to program. The possibility of a new AI being developed that is alien to mankind seems like the perfect program for world domination but is negated when considering that no one can create that which is beyond their capabilities. However, man is still capable of evil, and there is definitely a

concern that human evil will be programmed into an AI system, and should that appear in the healthcare industry, it may have devastating effects. One solution to this problem would be keeping the AI accountable to humans in all aspects. In medical facilities, AI would not operate independently of physicians but rather alongside doctors and nurses. Another solution similar to the first is creating committees to review these kinds of problems similar to a Hospital Ethics Committee, which reviews ethical dilemmas that appear during clinical practices (Hajibabaee). In case an AI provides a controversial solution or appears to act unethically, the program and decision would be brought to the review board to review its method, decision, pattern of decisions, and anything else to indicate a potential for harm. This would keep the AI accountable to humans and prevent it from acting free of oversight. Overall, artificial intelligence technology will not evolve beyond simply becoming advanced human-like beings because that is all humans know to program, and absolute oversight even to the level of a committee will help prevent AI from getting out of control.

The Hippocratic Oath states that a doctor should do no harm (Shiel). Unfortunately, medical errors compromise this oath far too often, but the development of artificial intelligence could help reduce these kinds of errors. Artificial intelligence is one of the most expansive and innovative technologies in the world, and with it, various industries like business and automotives are creating and implementing a new generation of technology that acts as a secondary mind with varying capabilities. Humans are not perfect, and mistakes in an industry dealing with peoples' lives are deadly. Preventing and reducing these kinds of mistakes all over medical facilities, whether in medicinal dosages in hospitals or general patient services in the waiting room, will help decrease the inordinately high number of people who die from medical errors every year and fulfill the Hippocratic Oath.

Works Cited

"The Anchoring Effect and How It Can Impact Your Negotiation." *Program On Negotiation*, Harvard Law School, 20 Sept. 2019, www.pon.harvard.edu/daily/negotiation-skills-daily/the-drawbacks-of-goals/.

Breunig, Matthias, et al. "Building Smarter Cars with Smarter Factories: How AI Will Change the Auto Business." *McKinsey Digital*, McKinsey & Company, Oct. 2017, www.mckinsey.com/business-functions/mckinsey-digital/our-insights/building-smarter-cars.

Carver, Niki, and John E Hipskind. "Medical Error." *NCBI*, U.S. National Library of Medicine, 15 May 2019, www.ncbi.nlm.nih.gov/books/NBK430763/.

Davenport, Thomas H., and Rajeev Ronanki. "3 Things AI Can Already Do for Your Company." *Harvard Business Review*, Harvard Business Publishing, 24 July 2018, hbr.org/2018/01/artificial-intelligence-for-the-real-world.

Dobelli, Rolf. "2015 : What Do You Think about Machines That Think?" *Edge.org*, Edge Foundation, Inc, 2015, www.edge.org/response-detail/26170.

Flouris, Andreas D., and Jack Duffy. "Applications of Artificial Intelligence Systems in the Analysis of Epidemiological Data." *European Journal of Epidemiology*, vol. 21, no. 3, 2006, pp. 167–170. *JSTOR*, http://www.jstor.org/stable/20445657

Gibbs, Samuel. "AlphaZero AI Beats Champion Chess Program after Teaching Itself in Four Hours." *The Guardian*, Guardian News and Media, 7 Dec. 2017,

www.theguardian.com/technology/2017/dec/07/alphazero-google-deepmind-ai-beats-champion-program-teaching-itself-to-play-four-hours.

Hajibabaee, Fatemeh et al. "Hospital/clinical ethics committees' notion: an overview." *Journal of medical ethics and history of medicine* vol. 9 17. 18 Dec. 2016

Heath, Sara. "Better Practice Operations Vital to Improve Patient Satisfaction." *Patient Engagement HIT*, Xtelligent Healthcare Media, 1 Sept. 2017, patientengagementhit.com/news/better-practice-operations-vital-to-improve-patient-satisfaction.

Kennedy, Lauren Paige. "How Artificial Intelligence Helps in Health Care." *WebMD*, WebMD, 29 Nov. 2018, www.webmd.com/a-to-z-guides/features/artificial-intelligence-helps-health-care#1.

McMains, Vanessa, and Lauren Nelson. "Study Suggests Medical Errors Now Third Leading Cause of Death in the U.S. - 05/03/2016." *Johns Hopkins Medicine*, The Johns Hopkins University, The Johns Hopkins Hospital, and Johns Hopkins Health System, 3 May 2016, www.hopkinsmedicine.org/news/media/releases/study_suggests_medical_errors_now_third_leading_cause_of_death_in_the_us.

Mnih, Volodymyr, et al. "Playing Atari with Deep Reinforcement Learning." *DeepMind Technologies*, 1 Jan. 2013, pp. 1–9., deepmind.com/research/publications/playing-atari-deep-reinforcement-learning.

Rossi, Francesca. "Building Trust in Artificial Intelligence." *Journal of International Affairs*, vol. 72, no. 1, 2019, pp. 127–134. *JSTOR*, http://www.jstor.org/stable/26588348.

Schroer, Alyssa. "Artificial Intelligence in Cars Powers an AI Revolution in the Auto Industry." *Built In*, 31 Jan. 2019, builtin.com/artificial-intelligence/artificial-intelligence-automotive-industry.

Shaw, Jonathan. "Artificial Intelligence and Ethics." *Harvard Magazine*, 30 Apr. 2019, harvardmagazine.com/2019/01/artificial-intelligence-limitations.

Shiel, William C. "Medical Definition of Hippocratic Oath." *MedicineNet*, MedicineNet, 6 Mar. 2018, www.medicinenet.com/script/main/art.asp?articlekey=20909.

Sraders, Anne. "What Is Artificial Intelligence? Examples and News in 2019." *TheStreet*, 3 Jan. 2019, www.thestreet.com/technology/what-is-artificial-intelligence-14822076.

"When Was the First Computer Invented?" *Computer Hope*, 2 Aug. 2019, www.computerhope.com/issues/ch000984.htm.

Wilson, Elizabeth A. *Affect and Artificial Intelligence.* University of Washington Press, 2010, pp. 32-182. *JSTOR*, http://www.jstor/stable/j.ctvcwnbsb.6.