The ethical implications of AI in healthcare

Ryan Scott

Mellon Research Fellowship, Brightpoint Community College

Midlothian, VA

rscott@brightpoint.edu

The rapid advancement of artificial intelligence (AI) in the last few years has had it encroaching in just about every field imaginable, including healthcare. In this paper, we will explore the ethical implications of AI in healthcare, focusing on its current use, ethical questions raised, incorporation into medical education, and the future clinical environment. By examining peer-reviewed research and incorporating perspectives from medical professionals, this paper aims to provide a comprehensive analysis of the ethical implications of AI in healthcare.

AI is already in use across many areas of medicine, including monitoring of vitals, diagnostic imaging such as autonomously spotting cancer cells [7], analyzing mental health risks [11] and even planning treatment for optimal outcomes. Data from a large survey conducted by the American Medical Association (AMA) suggests approximately one third of doctors are currently using some form of AI tools in their daily practice, albeit mostly for administrative tasks such as "…generation of patient messages and chart summaries, and prediction of demand and associated workforce needs." [13]. By utilizing a combination of machine learning algorithms, and modern computing, AI systems can analyze vast amounts of medical data at speeds impossible to manually evaluate, or even evaluate with mathematical modeling. After doing that, AI can in some cases provide accurate predictions and suggestions. However, this raises ethical questions that need to be addressed. Even if an AI driven medical system statistically outperforms our current healthcare system, it is unlikely to be perfect. What do we do if it makes a bad call? What is the patient perspective and handing control of their care to an algorithm? This is a question that requires more context to answer.

A particularly famous case of a bad ethical call related to healthcare was the case of Henrietta Lacks, an African American woman whose cancer cells led to the creation the HeLa cell line [14]. Unbeknownst to her, her “immortal” cells were passed along to researchers and used en masse. Ultimately this helped over 10 million people, but was it the right thing to do [14]? Numerically speaking, more people were helped than harmed but should we even be trying to quantify this kind of moral dilemma? The use of AI in healthcare raises several ethical questions in a similar vein. Jiang states “Before AI systems can be deployed in healthcare applications, they need to be ‘trained’ through data that are generated from clinical activities, such as screening, diagnosis, treatment assignment and so on, so that they can learn similar groups of subjects, associations between subject features and outcomes of interest.” [10]. When considering the methods of machine learning, is medical data collected from a patient able to be added to a model for use in the algorithm? Research from the University of Chicago suggests patients own their data “and can take their data from one provider to another,”. [2] There are privacy concerns “...inferences that an AI algorithm can make…by comparing patterns in the individual behavior to patterns seen across similar populations of individuals.”[2]. Besides the obvious questions regarding privacy (HIPPA was not enacted until 1996), there are “...inferences that an AI algorithm can make about the consumer by comparing patterns in the individual behavior to patterns seen across similar populations of individuals.” [2] There can be issues related to data ownership, informed consent, and bias in algorithmic decision-making. As a counterpoint, what is the difference between a learning model holding this data, and a medical professional? If people are not perfect, should we expect AI to be? This is in no way excusing the malicious actions of humans, but can we as non-perfect individuals expect to

create a perfect system? AI is a tool; it inherently has no motives as it is not a sentient being. However, a large model could be loaded with biased data, for example from only one group. For example, autoimmune issues are almost exclusive to first world countries. If we train a model saying that autoimmune diseases are a possible cause of X Y and Z symptoms, it is not prepared to be deployed in less developed countries. Is this omission negligent?

The non-perfect medical system starts with higher education, and then medical school. As the use of AI becomes more prominent in healthcare, there is good reason to evaluate the incorporation of AI into medical education. While AI has seemingly burst into the public eye only within the last few years, AI is not new news. Research from 1991 claims "...artificial Intelligence (AI) is having a large effect on the economy." [1]. Development of AI has been a long and steady process driven by advancements in software and hardware. Across science, medicine, education and the humanities, AI has had a significant impact since it was first developed. The education system as a whole has had to evolve to adapt with technology through several major leaps. Some major milestones we could consider would be trade based economies, the industrial revolution, the computing age, and arguably the current evolution of AI.

AI has changed the way people can learn, process and access information. Currently public perception of AI in education tends to be negative as it may inhibit the learning process by allowing students to take shortcuts. However, the question has to be raised, if these shortcuts will become commonplace in the workforce, are they shortcuts? Students no longer need to travel to a library to access encyclopedias or newspapers on microfiche. The internet

age has almost limitless information available with a few clicks. Data from Lee suggests allowing students quick and easy access to information with no social pressure was beneficial in a class content review environment [6]. Analogies to scientific calculators or spellcheck may seem apt but they are missing the same important point in degrading problem solving and critical thinking. Calculators, spellcheck, and Google were deemed bad for learning when they were the latest and greatest technology. However, in higher education the student still had to understand how to solve the problem to a degree. In the context of education, it is important to draw the line between information and a solution. While this is inherently vague, we can say information is like the pieces of a puzzle, and a solution is the put together puzzle. While both give you the pieces, the student needs to be able to put the pieces together to learn. Once the puzzle is completed, it literally leads to the complete picture, which helps reinforce the learning process. Metaphorically speaking, AI learning models get to look at so many pieces and so many puzzles that they can assemble the information on our behalf in many cases. This is hypothetical as we are looking to the future but to summarize, access to information is good to a point, that point being when the library of information is so vast and easily searchable it negates the need for students to critically think and analyze the problems they are presented with.

In a world where this vast easily searchable library of information exists, it inherently degrades the learning process if the student allows it to. Exams provide a roadblock for this, stopping many of these students in their tracks, which is a necessary evil. Circling back to Lee's study "Research results showed that the application of AI-based chatbots in the review process of public health courses could improve students' academic performance, self-efficacy, learning

attitude, and motivation" [6]. However, overusing a tool to the point it inhibits learning is a responsibility that lies with the student. This is especially important with the education of medical students, as they will have a higher ethical responsibility than most in practice. The education system may rightfully respond by weighing tests and labs more heavily since AI would not be available during, for example, a calculus exam or a chemistry lab. We do not need to tear up the education system and allow students to use AI on an exam, a reasonable and proportional response will suffice and occur naturally as it has for prior technological evolutions. Society requires more skilled and qualified doctors. Lowering standards of care is not the solution here. This is especially important in medical school. Is it ethical to allow medical students to rely on an AI model to get through their education? Is it fair to the patient to allow those same students to practice medicine?

AI in the business world has been hailed as a time saving discovery, healthcare may not be a business, but in terms of work hours we can make a comparison. It is undeniable that the world needs doctors, and more of them. The World Health Organization (WHO) indicates there is a global deficit of 4.3 million physicians, nurses, and health professionals. [15] While lowering standards is not the solution, people who once had to work on repetitive tasks can benefit from AI as machine learning can manage the task and the person can focus on other parts of their job the business finds more valuable. Similarly, repetitive tasks in healthcare can be automated, allowing the same number of doctors to provide the same quality of care to more patients. Using AI can eliminate human data entry errors. With medical malpractice being one of the leading causes of death in the U.S. we cannot idealize the current healthcare system to the point where we view it as unchangeable [18]. Extensive literature research summarized it very

well, referring to the "...dual advantage of ethical AI; maximizing the opportunities to cut costs, improve care, and improve the efficiency of health and care systems, whilst proactively avoiding the potential harms." [12].

Looking at medicine from a diagnostic or analytical point of view, we can start to see some data that statistically suggests we should begin to meld AI and diagnostics. Looking at data from Morley's study regarding early detection of strokes "The data would be extracted and modeled by hidden Markov models and SVM, and the algorithm could correctly classify 90.5% of the subjects to the right group." [12]. This is breakthrough levels of diagnostic improvement, currently strokes are primarily diagnosed after symptoms show. Is it ethical to withhold this technology due to red tape?

That is a dangerous question, but it is important to think about. If technology is evolving faster than the regulatory bodies of medicine can handle, what do we do? This is where it becomes an ethical question instead of an objective question, and it is quite possible that there is no correct answer**.** Despite the seemingly limitless benefits, it will be vital to critically analyze the role and limitations of AI in medicine to ensure ethical and responsible implementation. Figure 1 below from the article by Kun-Hsung summarizes the possible progression of clinical practice with implementation of AI.

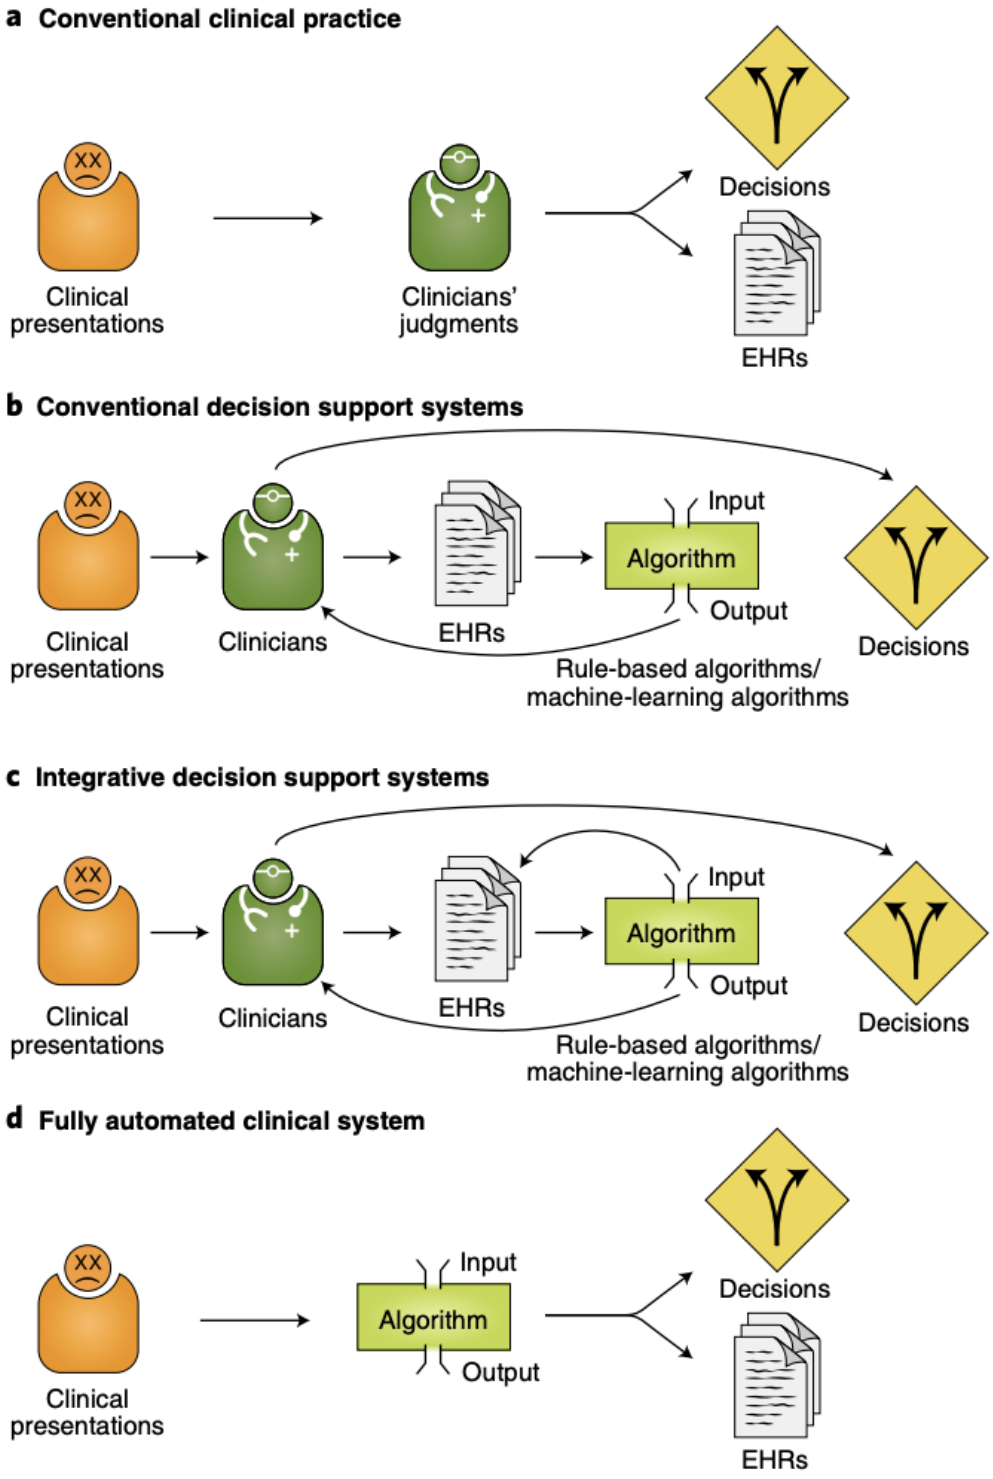

Figure 1a-1d: possible implementations of ML tools in the clinical environment [3]

Starting with traditional clinical practice shown in figure 1a, when a patient arrives to see the clinician, the clinician uses their judgment and knowledge to assess the patient and make decisions. Electronic Health Records (EHR) are generated manually, which is a huge work obligation and often runs over into the medical professional's personal life. This is the system we use currently, and there is certainly the potential for ethical issues as we have discussed previously. Should we expect an AI implementation to be both statistically superior, and more ethical than the humans whose decisions are used to train the algorithm?

Adding an AI tool shown in figure 1b, after a clinician examines the patient and then the EHR is sent to the ML algorithm. The clinician then uses the output of the model to facilitate making a diagnosis. While it does not pertain to the ethics, this may actually make the clinician's job more difficult as the EHR's need to be completed on the fly, which is simply unrealistic oftentimes. Ignoring the potential workload issue, this seems like a reasonable method of logically validating the clinician's diagnosis, but it certainly leaves room for the algorithm influencing the decision making. However, is the influence a bad thing if a doctor is ultimately making the call? The answer is it depends. Does the doctor blindly follow what the ML model says? Or do they work in parallel so the conclusions can be compared? If we assume the doctor maintains their professional and ethical standards, conceptually there should not be an issue with this method of clinical integration.

In the model shown in figure 1c, additional information is added to the patient's EHR after being run through the ML algorithm. This is very similar to the model shown in figure 1b. Having the entire objective story in EHR would be a positive because data entry errors are a

significant and documented problem in many fields. With that being said, do we want to allow AI to do write to EHR's unchecked? As a counter point, if it is more accurate than a person, is it ethical to not implement automated EHR writing for the benefit of the patient?

Fully automating as shown in figure 1d removes a clinician from the practice altogether and upon first glance, can be ethically shaky at best. So many conditions present the same way and not having a clinician in the loop could result in inaccurate diagnoses, which is obviously a negative for the patient. As an example of a fully automated diagnostic tool, at Duke University, an app for parents to screen their children for autism spectrum disorder (ASD) is being trialed. The app uses a phone camera "to record young children's reactions while they watch movies designed to elicit autism risk behaviors, such as patterns of emotion and attention, on the device's screen." [7]. This use of fully automated data collection would be helpful as bringing a young child in for lengthy observation is not always practical. The ethics of data protection are being addressed as researchers "develop new machine learning algorithms for privacy filters for the images and videos they collect." [7]. This is a huge leap from the model suggested in figure 1c, and while we may have evidence to suggest it can be helpful in certain situations. I propose a fifth model may fill a gap that could provoke some interesting questions.

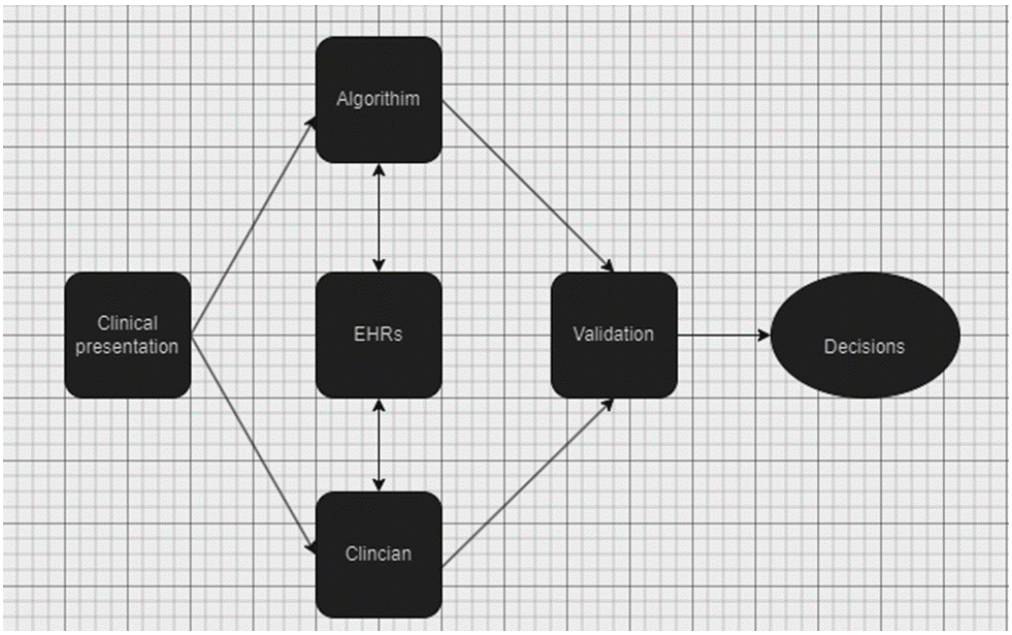

In this conceptual model, the doctor examines the patient with the ML tool listening to the exam and recording the relevant points of the conversation. Additionally, test readings can be automatically written to the EHR. The AI tool would also be able to review the patients' EHR at the same time for any relevant information. The EHR could then be validated by the medical professional, and the decision could be made with the doctor and AI tool in parallel. This significantly could reduce the workload of the doctor regarding EHR's, but maintain a human validation at every step of the process.

As the above example suggests, it will be vital to critically analyze the role and limitations of AI in healthcare to ensure ethical and responsible implementation. There may not be a correct answer at this point, and there is most certainly some opinion involved. The opinions of both the patient and medical professional are relevant.

We need to examine what to do and who is responsible in the event of medical mistakes driven by AI tools before we can proceed further with this concept. Recently a chatbot provided incorrect information to a customer of Air Canada regarding a flight [16]. Imagine a hypothetical scenario where an automated medical system gave an incorrect recommendation of a medicine the person is allergic to. In the event of accidents or malpractice caused by AI systems, it is important to determine who is responsible to analyze significant ethical and legal concerns.

Discussions surrounding liability issues and appropriate measures to hold individuals or organizations accountable need to be put in place before launching AI powered patient information systems without a human in the loop. In an objective and emotion-free world, we would fully integrate any tested and proven AI tool into healthcare that outperforms humans. However, considering the great potential of the human psyche. In the ultimate worst-case scenario, we cannot expect people to throw up their hands and say oh well if someone is killed as a result of malpractice originating from an AI healthcare tool, it cannot be the fault of the medical professional. However, is it not their responsibility? To make an argument in favor of allowing AI to seep into healthcare, currently the doctor is responsible for the care of the patient. That remains true regardless of the tools they use. We can look forward and apply this reasoning to medical professionals using AI powered tools, if and only if a qualified medical professional stays in the loop. Allowing AI the autonomy to make decisions in a vacuum is an entirely different discussion. Maintaining the ethical order that has been established in medicine requires a trained and qualified medical practitioner to be involved in the decision making process as they currently are.

AI systems have the potential to revolutionize healthcare by automating routine tasks and enhancing diagnostic processes. Healthcare systems and providers should not fully rely on AI systems and ensure human involvement in critical aspects of healthcare. Patients want to have relationships with their medical providers and be able to trust their diagnoses and recommendations. Olaisen's research shows "The quality of the physician-patient relationship is positively associated with functional health." [17]**.** Having AI in place behind the scenes analyzing and processing data can allow medical providers to make more informed decisions. The now common phrase "human in the loop" is especially important in the clinical environment. To summarize, the near future and far future clinical environments entails two different questions. With the current intelligence level of AI models, the studies show it would be beneficial to integrate AI into the medical process to analyze data, as long as a physician makes the ultimate call for diagnoses and treatment.

Using AI in healthcare presents both incredible opportunities and ethical dilemmas. If generalized artificial intelligence is realized, we are drawn to some questions. If the technology we create is more advanced than we are, who is in the driver's seat? Pertaining to healthcare, currently we feel obligated to provide patients with the best care possible. Do we continue to do that even if it involves using tools and systems entirely run or developed by AI? Ultimately, we cannot conclude that there is one correct answer. If implementing AI is the statistically right choice in that more patients have better outcomes, should we? Likewise, is this a reduction in our autonomy and self-sufficiency as a society? Even if that is the case, can we say it is right to willfully have worse healthcare outcomes in the name of autonomy? To reflect on the bigger picture, are we obligated to become as advanced of a society as possible?

References

This journal article has provided me with an interesting perspective on how AI was viewed in the 1990s. The author was one of the leading experts in computer science and AI at the time this article was written. I will use this information in combination with the present general information articles to help form the background section of my paper. Additionally, I can compare predictions of the past to the present outcomes and use that to gain a different perspective on present predictions.

[1] R. C. Schank, "Where's the AI?," AI Magazine, vol. 12, no. 4, p. 38, 1991.

This journal article takes an analytical perspective on the impact of AI on the economy and then discusses potential issues from a more theoretical perspective. Direct examples of human impact are discussed with the concern of AI taking over jobs and raising the skill floor. There are many technical economic terms used, and I may need to include an additional source to provide background. This will be a strong article to use in both my background and my discussion.

[2] J. Furman and R. Seamans, "AI and the Economy," Innovation Policy and the Economy, vol. 19, 2019.

This journal article discusses experimental data from using AI to analyze data in healthcare. Additionally, it covers some of the potential social, economic, and legal challenges. The three authors are either medical professionals, computer scientists, or both. This article is highly

credible with a large quantity of experimental data and a solid discussion. This will be a primary source in my paper's background and analysis sections.

[3] Y. Kun-Hsung, B. Andrew L. and K. Issac S., "Artificial intelligence in healthcare," Nature, no. 2, pp. 719-731, 2018.

This web article provides me with background information about the issues of AI and intellectual property (IP). The author is a highly experienced technology writer, but not an expert in the field itself. This article will primarily be used to cite the issues with AI and IP. It will be backed up by a more credible primary source.

[4] R. Rao, "spectrum.ieee," IEEE Spectrum, 13 June 2023. [Online]. Available: https://spectrum.ieee.org/generative-ai-ip-problem. [Accessed 2 October 2023].

This journal article discusses the moral, ethical, and legal implications of “accident algorithms.” In other words, what should something like a self-driving car do when it could save its occupants at the expense of another life? The author is a professor of ethics and artificial intelligence at LMU Munich, and the article is credible for the scope it covers. This article talks about the problem rather than proposing solutions, I think it will contribute more to my background than my discussion.

[5] S. Nyholm and J. Smids. “The Ethics of Accident-Alogritims for Self-Driving Cars: an Applied Trolley Problem?, ”Ethical Theory and Moral Practice, vol 19, no. 5, pp. 1275-1289, 2016.

This journal article talks about the potential positive impacts of integrating AI into the college level education system. The author is credible having published multiple studies in education technology research. This article reviews a case study which they performed that showing AI when implemented correctly can significantly help students' learning outcomes. This provides a strong counterpoint to the popular opinion that AI generally inhibits learning.

[6] Y. Lee, G. Hwang, and P. Chen, "Impacts of an AI-based chatbot on college students' after class review, academic performance, self-efficacy, learning attitude, and motivation," Educational technology research and development, vol. 70, pp. 1834-1865, 2022.

This magazine article provides me with popular opinion regarding the use of AI in healthcare. The magazine is published by Duke University ECE department and is credible for this source. This article will be primarily used in the background section to establish the existing perception of AI in healthcare.

[7] "AI in health," Duke ECE, pp. 1-20, 2020

This journal article provides me with arguments for and against the regulation of AI. The author is a professor of law and economics at an established law school. This article is credible. I hope to be able to use the regulatory perspective in collaboration with the more specific information provided by other sources.

[8] M. C. Buiten, "Towards Intelligent Regulation of Artificial Intelligence," European Journal of Risk Regulation, vol. 10, no. 1, pp. 41-59, 2019.

This source provides me with an interesting look into the past regarding AI and autonomous vehicles. The corporate author is a research division of the US military, which is credible in this context. I am not sure how this source will be used in my paper, but it is interesting enough that I want to include it if possible.

[9] Defense Advanced Research Projects Agency, "darpa," 2005. [Online]. Available: https://www.darpa.mil/about-us/timeline/-grand-challenge-for-autonomous-vehicles. [Accessed 11 November 2023].

This peer reviewed journal article provides me with a high-level survey of the uses of AI in healthcare. This is a highly technical article that looks at things from more of a research perspective. The author is highly credible, with numerous publications and years of medical experience. This article has been cited nearly 3000 times. I think this article will be strong in my analysis section because it talks about what AI is actually being used for and how it works, as opposed to just a study of one specific topic.

[10] F. Jiang et al., “Artificial Intelligence in healthcare: Past, present and future,” Stroke and Vascular Neurology, vol. 2, no. 4, pp. 230–243, Jun. 2017. doi:10.1136/svn-2017-000101

This article, which is a documentation of a summit regarding AI in healthcare talks about some more specific implementations of AI and healthcare and goes into some interesting examples of clinical use. The author is credible, with over 30 peer-reviewed publications and a career relevant to the field of discussion. I am not sure exactly how I will use this article in my paper

yet, but I think the examples of clinical applications will really help me analyze this from a humanities perspective and tie things together.

[11] E. Drysdale, “Implementing AI in healthcare,” The Hospital for Sick Children and the Vector Institute for Artificial Intelligence, Mar. 2020.

This a gem of an article for my research topic. It reviews the incentives for introducing AI into healthcare, while also discussing the potential ethical issues. The author is highly credible. The article has been cited several hundred times. I think this article will help me connect the raw data from my other sources together into a coherent research argument.

[12] J. Morley et al., “The ethics of AI in health care: A mapping review,” Social Science &; Medicine, vol. 260, p. 113172, Sep. 2020. doi:10.1016/j.socscimed.2020.113172

This survey conducted by the American Medical Association (AMA) gives me a data set with over 1000 samples that I can use to understand the general sentiment of AI in healthcare from the perspective of the medical community. The author is not an expert in the field but the publisher being the AMA provides this article with a high degree of credibility regarding the data itself.

[13] A. Robeznieks, Ed., “AMA Augmented Intelligence Research,” Big majority of doctors see upsides to using health care AI, https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf (accessed Feb. 1, 2024).

This book provides me with background information for an instance of unethical medical practice. While the book itself is not peer reviewed, it cites multiple peer reviewed sources and

was a New York Times Bestseller. I use the information from this book to make an example of poor ethical decision making in healthcare, and what that can do to people.

[14] R. Skloot, “The Immortal Life of Henrietta Lacks,” New York City, Crown Publishers, 2010.

This peer reviewed journal article provides me with detailed documentation of the global healthcare worker shortage. The author is credible with multiple research papers on this specific topic, which pertain to his employment in research at a major university.

[15] C. Aluttis, T. Bishaw, and M. W. Frank, “The workforce for Health in a globalized context – global shortages and international migration,” Global Health Action, vol. 7, no. 1, Feb. 2014. doi:10.3402/gha.v7.23611

This web article was in my news feed and I thought I may be able to make an interesting connection to healthcare given that the airline was held liable. I plan to use this as an ethical example which I can project towards healthcare and ask a “what if”question. Specifically, I want to ask what if this type of error happened in a healthcare setting.

[16] M. Yagoda, “Airline held liable for its chatbot giving passenger bad advice - what this means for travellers,” BBC News, https://www.bbc.com/travel/article/20240222-air-canada-chatbot-misinformation-what-travellers-should-know (accessed Mar. 15, 2024).

This article provides me with documented statistical analysis of outcomes and the doctor patient relationship. The study found that having a good relationship with a doctor is correlation to improved patient outcomes. I plan to use this to argue for maintaining human contact for patients in clinical settings.

[17] R. H. Olaisen et al., “Assessing the longitudinal impact of physician-patient relationship on Functional Health,” The Annals of Family Medicine, vol. 18, no. 5, pp. 422–429, Sep. 2020. doi:10.1370/afm.2554

This article is used exclusively to provide me with a documented statistic regarding the number of fatalities caused in the US by medical malpractice. This is used to inform the reader that the current medical system, while it has progressed over time, is not perfect.

[18] R. Sipherd, “The third-leading cause of death in us most doctors don’t want you to know about,” CNBC, https://www.cnbc.com/2018/02/22/medical-errors-third-leading-cause-of-death-in-america.html#:~:text=A%20recent%20Johns%20Hopkins%20study,after%20heart%20disease%20and%20cancer. (accessed Apr. 26, 2024).

Appendix A

During my time working on this project, I was presented with an opportunity to present my work at the Virginia Humanities Conference (VHC). My research won the best undergraduate presentation award and I have included the award statement.

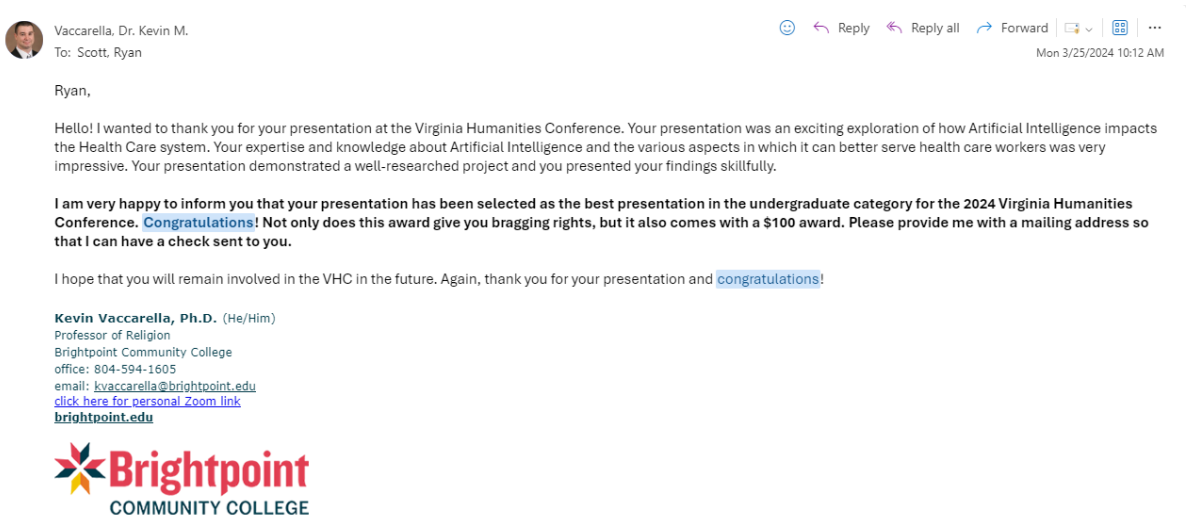

Vaccarella, Dr. Kevin M.
To: Scott, Ryan

Reply Reply all Forward
Mon 3/25/2024 10:12 AM

Ryan,

Hello! I wanted to thank you for your presentation at the Virginia Humanities Conference. Your presentation was an exciting exploration of how Artificial Intelligence impacts the Health Care system. Your expertise and knowledge about Artificial Intelligence and the various aspects in which it can better serve health care workers was very impressive. Your presentation demonstrated a well-researched project and you presented your findings skillfully.

I am very happy to inform you that your presentation has been selected as the best presentation in the undergraduate category for the 2024 Virginia Humanities Conference. Congratulations! Not only does this award give you bragging rights, but it also comes with a $100 award. Please provide me with a mailing address so that I can have a check sent to you.

I hope that you will remain involved in the VHC in the future. Again, thank you for your presentation and congratulations!

Kevin Vaccarella, Ph.D. (He/Him)
Professor of Religion
Brightpoint Community College
office: 804-594-1605
email: kvaccarella@brightpoint.edu
click here for personal Zoom link
brightpoint.edu

Brightpoint
COMMUNITY COLLEGE