The Ethical Analysis of the Integration of Machine Learning within Patient Diagnostic Imaging Procedures

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Nyota Patel

Spring 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor
Dr. MC Forelle, Department of Engineering and Society

Introduction

A common perception held by many individuals is that artificial intelligence (AI) has the potential to transform a wide variety of industries by improving worker productivity and optimizing workflow (Nadimpalli, 2017). However, Elon Musk, one of the most successful entrepreneurs of all time, had the following statement to say about artificial intelligence: “mark my words, AI is far more dangerous than nukes” (Clifford, 2018). Despite apprehensions such as these, AI is growing rapidly and is being implemented in several fields including business, transportation, and even healthcare to aid in the decision-making process, increase workflow efficiency, and reduce costs (Nadimpalli, 2017). Despite these benefits, this technology has many risks such as loss of employment, reduction in disposable income, and health issues brought upon by biases within the AI algorithms (Nadimpalli, 2017).

In healthcare, specifically, machine learning (ML) algorithms are being used to diagnose patients, monitor health, develop drugs, and manage medical data (Amisha et al., 2019). ML is a branch of AI that allows computers to learn from a set of data on how to perform a specific task without the pre-defined rules (Rajkomar et al., 2018). The integration of ML within healthcare raises several concerns ranging from privacy to ethical issues. In order to perform their intended function, ML algorithms require a large amount of patient information. In the event of a cyberattack, patient information stored within these algorithms can be compromised, leading to potential data breaches and identity theft. Additionally, the use of these algorithms within medicine can raise ethical dilemmas on who should be held accountable in the event of an error. Some individuals may argue that the software engineers should be held responsible for this mistake because they created a biased algorithm. Others, however, may argue that the physician should be held liable as they are the ones who are employing the use of this technology.

It is imperative to address issues such as these before adopting widespread use of AI. As mentioned before, artificial intelligence is a rapidly evolving field with many applications in a wide variety of industries. It is of crucial importance for researchers to explore the benefits and risks associated with this technology to prevent security risks and ensure ethical principles are being adhered to. Additionally, researching the advantages and challenges associated with AI can ensure these algorithms are being used and integrated in a secure and responsible manner into various fields, particularly healthcare. Lastly, this work holds significant importance because it can help guide the Food and Drug Administration (FDA) in creating regulatory policies that ensure the integration of ML into healthcare upholds medical ethics.

In this paper, I argue that while the implementation of ML-based algorithms within patient diagnostic imaging procedures improves diagnostic accuracy, it ultimately violates the four pillars of medical ethics. In order to support my argument, I will first provide a literature review on the applications of machine learning algorithms within healthcare and their benefits in regard to increasing diagnostic accuracy and improving clinical efficiency. I will then present information about the origins of the four pillars of medical ethics and how ML impacts these four principles. Then, I will analyze data from different medical studies and academic papers to conduct an ethical analysis. The data that I will analyze will be authored by several individuals with diverse backgrounds ranging from engineers to doctors to researchers. Through this analysis, I find that using ML-based algorithms in patient diagnostic imaging procedures violates the autonomy, non-maleficence, and justice principles while upholding the beneficence principle. Finally, I will end this paper with a discussion about the limitations of my research and how my work can be used by the FDA to develop a pathway for AI and ML-based technologies to obtain premarket approval.

Literature Review

A wide variety of studies have been conducted to determine the accuracy and efficiency of machine learning algorithms in diagnosing patients within healthcare. According to a study conducted at the Johns Hopkins University, approximately 40,500 patients die each year due to diagnostic errors. Because of these errors, health care expenses increase by approximately $300,000 for every malpractice claim that is filed for a misdiagnosis (Dilsizian & Siegel, 2014). The use of AI in accordance with radiologists can help mitigate these challenges. For instance, a study conducted in the radiology department of a cancer center found that the collaboration between radiologists and AI not only increased the diagnostic accuracy of detecting incidental pulmonary embolism (IPE) on CT scans but it also significantly reduced the missed rate of IPE from 44.8% to 2.6% (Topff et al., 2023). Additionally, Figure 2 shows how AI reduced the median detection and notification time for IPE from several days to 1.0 hour (Topff et al., 2023).

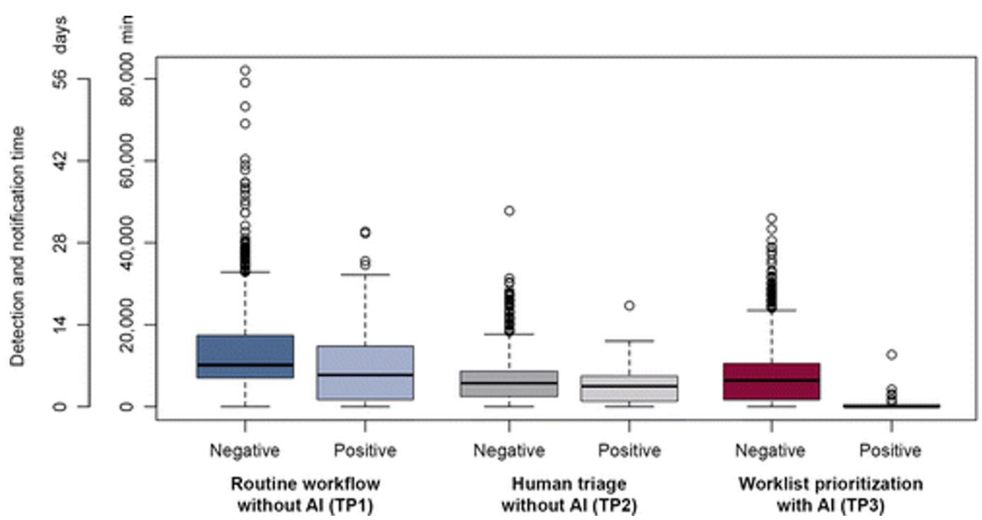

Figure 1: Box plot comparison of the detection and notification time of IPE–positive and –negative CT scans with and without the use of AI (Topff et al., 2023).

The accuracy of these AI algorithms can extend to other illnesses and imaging techniques. For example, in a study conducted by Nishida et al. (2022), three ML models were developed using an ultrasound image dataset to diagnose liver tumors. A comparative analysis was then performed to differentiate the diagnostic accuracy of the ML models against that of physicians (Nishida et al., 2022). The study found that all three algorithms were able to diagnose the liver tumors at an accuracy that was significantly higher than that of human physicians (Nishida et al., 2022). Beyond achieving a high diagnostic accuracy, ML can also help improve clinical workflow by scheduling appointments, digitizing medical records, and increasing patient-physician interaction (Amisha et al., 2019). Although it is evident that extensive research has been conducted on the benefits of using ML-based algorithms within healthcare, the impact of ML on the medical ethical framework needs to be researched further before this technology is integrated into medicine.

The medical ethical framework first originated in the 5th century B.C. in Greece with the introduction of the Hippocratic Oath (Veatch, 1997). This Oath required all healthcare professionals to swear upon Greek gods to abide by the ethical principles and act in the patient's best interests (Veatch, 1997). Based off this Oath and previous works such as Percival's *Medical Ethics*, the American Medical Association (AMA) published a document, entitled the *Code of Medical Ethics*, in 1847 (Veatch, 1997). Over time, both the Hippocratic Oath and the *Code of Medical Ethics* underwent several revisions to emphasize patients and inclusivity of all religious beliefs (Veatch, 1997). From this Oath and previous documents, in 1979, two American philosophers, Beachump and Childress, published a book called the 'Principles of Biomedical Ethics,' which presented the four pillars of medical ethics (Aksoy & Tenik, 2002; Varkey, 2021).

The foundation of modern medical ethics rests upon the following four principles: autonomy, beneficence, non-maleficence, and justice. Autonomy describes the right of the patient to make rational decisions and moral choices regarding their health (Varkey, 2021). The principle of beneficence obligates physicians to act in the patient's best interests, while the non-maleficence principle requires healthcare professionals to impose no harm to the patient (Varkey, 2021). Lastly, the principle of justice states that all patients should be treated fairly while ensuring they are given equitable access to healthcare resources (Varkey, 2021).

Integrating machine learning within patient diagnostic imaging procedures can have an impact on all four of these pillars. For instance, assume an algorithm was developed from a White male population and applied to diagnose a female population of all races (AlHasan, 2023). The nonmaleficence principle could be potentially violated due to a biased training dataset, making this algorithm prone to providing information that is either irrelevant or even harmful to the female population (Char et al., 2020). The autonomy principle can also be violated because a potential lack of communication between the doctors and engineers about the algorithm can make it difficult for both the patients and physicians to make informed decisions (Feudtner et al., 2018). The justice principle would also be violated because not all hospitals will have access to this technology. To ensure these principles are upheld, the following factors should be considered while designing, implementing, and validating the algorithm: communication, having a representative training dataset, and promoting equity in patient outcomes and resource allocation (Char et al., 2020; Rajkomar et al., 2018). However, further research still needs to be conducted to determine how the integration of machine learning within patient diagnostics procedures impacts medical ethics.

To help guide my analysis of how the four pillars of medical ethics are impacted by ML, I will use Trevor Pinch and Wiebe Bijker's theory, the Social Construction of Technology (SCOT) to identify the relevant social groups that are impacted by ML (Pinch & Bijker, 1984). According to SCOT, the creation of a technology is shaped by various factors, one of them being the relevant social groups (Pinch & Bijker, 1984). Relevant social groups can be defined as "all members of a certain social group [that] share the same set of meanings, attached to a specific artefact" (Pinch & Bijker, 1984, p. 414). In this paper, rather than applying SCOT in a conventional manner to determine how the relevant social groups influenced the development of ML algorithms, I will use SCOT to identify the relevant social groups that are impacted by ML. For this paper, the relevant social groups that I will focus on are software engineers, doctors, and patients. I will then determine how the integration of machine learning algorithms into patient imaging diagnostic procedures impacts these groups.

Methods

To answer my research question, I primarily gathered secondary sources from papers within academic journals regarding machine learning and medical ethics. The academic papers were written by individuals from various fields including medicine, research, and computer science. I then reviewed media and journalistic accounts about these groups to learn more about how ML healthcare applications can impact them. Additionally, I used the databases within UVA Library and reviewed the American Medical Association (AMA) *Code of Medical Ethics* to learn more about the ethical framework and how ML can impact this framework. For my primary source, I reviewed legal documents from the FDA to gain insights into the current regulations regarding ML in medicine. Then, with all this data, I conducted an ethical analysis to determine how medical ethics and the relevant social groups are impacted by the use of ML algorithms within patient diagnostic imaging procedures. In this analysis, I focused on both a quantitative and a qualitative approach, where the quantitative aspect included information about diagnostic errors and accuracy. The qualitative approach, on the other hand, focused on the social and ethical implications of ML-based healthcare algorithms.

Analysis

Integrating machine learning into patient diagnostic imaging procedures can help reduce diagnostic errors and improve patient outcomes, thus aligning with the beneficence principle of the four pillars of medical ethics. A convolutional neural network (subset of ML) algorithm called CheXNet was developed to diagnose pneumonia from chest X-rays, shown in Figure 1 (Rajpurkar et al., 2017). The performance of CheXNet in diagnosing pneumonia was statistically significantly higher than the radiologist's performance (Rajpurkar et al., 2017). This highlights that ML-based algorithms can not only detect patterns missed by physicians but also produce consistent results. Additionally, unlike most physicians, this algorithm is not influenced by external factors such as experience and emotion, which helps increase diagnostic accuracy.

Figure 2: CheXNet takes in the chest x-ray as the input and outputs the location of where pneumonia is the most severe, show in red at the lower left lobe (Rajpurkar et al., 2017).

In another study, a deep learning (subset of ML) algorithm's ability to predict the diagnosis of Alzheimer's Disease was compared to the performance of radiologists (Ding et al., 2019). This study demonstrated that the algorithm outperformed the radiologists by achieving an 82% specificity at 100% sensitivity, whereas the radiologists achieved a 57% sensitivity at 91% specificity (Ding et al., 2019). The algorithm's high sensitivity suggests early detections of Alzheimer's Disease can be made, which allows for earlier treatment and improved prognosis for these patients. Additionally, this algorithm can more consistently predict diagnoses of Alzheimer's Disease than human physicians. A doctor's performance is influenced by a variety of factors such as stress and fatigue. However, an algorithm's performance is based solely on previous datasets, which are used to train the model in order to accurately diagnose patients.

AlHasan (2023) and others have argued that if an unrepresentative dataset is used to train the model, the beneficence principle is violated due to the introduction of biases into the algorithm. However, most people fail to consider that this would violate the non-maleficence principle, not the beneficence principle. In order for the beneficence principle to be truly violated, the physician has to use the algorithm after being made aware of its biases. In cases such as these, the physician would not be acting in the patient's best interest, and the beneficence principle would be violated due to the physician's negligence, not because of the ML algorithm itself. Therefore, ML-based healthcare algorithms can increase diagnostic accuracy and improve patient outcomes without violating the beneficence principle. However, it is still important to consider how the non-maleficence principle is violated due to different sources of algorithm bias.

Biases in interactions with clinicians and in the model design violate the non-maleficence principle. One example of bias that occurs when interacting with physicians is automation bias (Rajkomar et al., 2018). This occurs when clinicians place too much trust in the algorithm without realizing that the model is less accurate for certain groups (Rajkomar et al., 2018). Another type of bias is label bias, which occurs within the model design. In this bias, a certain label does not hold the same meaning for all patients (Rajkomar et al., 2018). A blind reliance in these algorithms can lead to physicians making incorrect decisions that could potentially be detrimental to a patient's health. For instance, if an algorithm with automation or label bias was applied in a clinical setting, certain minority groups would be discriminated against and unable to reap the benefits of this algorithm. These biases within the algorithm can not only exacerbate social issues but also lead to misdiagnosis and inadequate treatment for certain marginalized groups, which would violate the non-maleficence principle of the four pillars of medical ethics.

The non-maleficence principle is also violated with the use of ML algorithms when an unrepresentative dataset is used. When developing a machine learning healthcare application, the algorithm can introduce racial bias, which can have harmful effects to a patient's health. Racial biases are introduced due to the inherent nature of health care varying by race (Char et al., 2018). Illnesses and treatments differ by race; thus, it is difficult to create an algorithm that can be generalized to all populations. The algorithm's bias can be amplified due to the underlying biases within the training data itself (Char et al., 2020). In order to accurately diagnose patients using ML algorithms, the algorithm first has to be trained, validated, and tested using patient datasets. If the dataset is biased, then the algorithm's bias is exacerbated by an unrepresentative training sample. Predictive scores generated from ML-based systems in regards to a patient's health have already failed due to a biased training set, which has led to racially discriminatory outcomes (Char et al., 2018). For instance, when data from the Framingham heart study was used to predict the likelihood of cardiovascular events in nonwhite populations, it led to biased results, with both overestimations and underestimations of risk because this study included data from a primarily white population (Gijsberts et al., 2015). Discriminatory outcomes such as these can not only be detrimental to the health of minority races but also amplify healthcare disparities.

The exacerbation of healthcare disparities due to algorithm bias violates the justice principle of the four pillars of medical ethics. For example, suppose an ML algorithm is developed to accurately diagnose acute coronary syndrome (ACS) in patients. In order to train, validate, and test this algorithm, data from the Framingham heart study is used. Since this algorithm lacks a representative dataset, it is unable to identify risk factors for minority populations and does not take into account the health conditions that are prevalent among nonwhite populations. Thus, when this algorithm is applied in a clinical setting, patient outcomes would only improve for the white population due to early diagnosis. However, for nonwhite populations, diagnostic errors would be more prevalent due to biases within the dataset, leading to inadequate or delayed treatment along with poor prognosis. This algorithm would, thus, worsen the existing healthcare inequalities as the morality rate from ACS is 30% higher among African Americans compared to non-Hispanic whites (Graham, 2015). In another case, an algorithm was developed to distribute healthcare resources and provide personalized care to patients depending on the severity of their illness (Ledford, 2019). This algorithm was widely used in American hospitals, and it was found that the algorithm was less likely to refer black patients for personalized care, despite them being equally sick as the white patients (Ledford, 2019). The algorithm assigned lower risk scores to black patients than equally sick white patients, most likely due to a biased training dataset (Ledford, 2019). Algorithms such as these amplify healthcare disparities and violate the justice principle because certain racial groups, particularly non-Hispanic white individuals, are allocated better resources and care while African American patients are more prone to receiving inadequate treatment.

In addition to racial discrimination, machine learning algorithms widen the socio-economic gap while violating the justice principle of the four pillars of medical ethics. ML algorithms are heavily dependent upon patient data from electronic health record (EHR) systems. Patient-reported outcomes “tend to be documented by individuals with higher income, younger age, and white race” (Gianfrancesco et al., 2018). Vulnerable populations such as individuals with a lower socioeconomic status and immigrants would have inconsistent or incomplete documentation due to receiving treatment from multiple institutions, unequal access to online patient, and lower health literacy (Gianfrancesco et al., 2018). Thus, if patient data from EHR was used to train an ML algorithm, it would produce inaccurate results and would be unable to adequately treat the health issues of these vulnerable populations. Furthermore, many of these algorithms are trained and tested in more privileged hospitals, who primarily serve higher income individuals (AlHasan, 2023). Many under-resources hospitals may not have access to this technology, and if they were to, these algorithms would not be applicable to a hospital that primarily serves lower-income individuals. Therefore, due to unequal access to healthcare resources, patients from marginalized populations are less likely to benefit from these algorithms, which further widens the socioeconomic gap and violates the justice principle.

Lastly, the autonomy principle of the four pillars of medical ethics is also violated when ML algorithms are developed without proper communication between the stakeholders. Physicians and patients are better able to make informed decisions regarding a patient's health when the goals and intentions of a machine learning application are clearly communicated. For instance, a study concluded that proper communication with stakeholders during the software development process can improve customer satisfaction and result in lower defect rates (Bakalova & Daneva, 2011; Korkala et al., 2006). This is likely because the stakeholders can provide consistent feedback on the algorithm, continuously assess whether the algorithm is meeting the objectives, and determine whether the algorithm's outputs align with the intended purposes of the algorithm (Char et al., 2020). This communication between the stakeholders helps increase the acceptance of the algorithm from the user's perspective. If, however, the stakeholders are not involved during the software development process, the physician and patient's ability to make rational decisions is compromised because they are not provided adequate information about the algorithm and its intended use. This lack of communication about the algorithm violates the autonomy principle and diminishes stakeholder trust and comfort in the algorithm. Therefore, it is imperative for stakeholders to follow this collaborative approach to ensure that physicians and patients are able to make rational decisions about a patient's health.

Conclusion

In this paper, I have argued that using ML-based algorithms within patient diagnostic imaging procedures upholds the beneficence principle due to an increase in diagnostic accuracy. However, ML healthcare applications ultimately violate the non-maleficence, justice, and autonomy principles while systematically discriminating against minority groups, exacerbating healthcare disparities, and widening the socio-economic gap. From this paper, the relevant social groups (software engineers, healthcare professionals, and patients) are able to gain a deeper understanding of how ML can impact the ethical framework, specifically the four pillars of medical ethics. Additionally, this analysis can provide valuable insights for regulatory bodies.

Currently, all medical devices are required to undergo the appropriate premarket pathway to be approved (U.S. Food and Drug Administration, 2024). For ML-based healthcare algorithms, there are three levels of clearance: 510(k), the de novo pathway, or premarket approval (Benjamens et al., 2020). The appropriate pathway for approval varies depending on the classification of the medical device. Since algorithms are continuously being updated, it is difficult for these technologies to go through the former approval process (Benjamens et al., 2020). Thus, currently, there is no formal approval process for ML-based healthcare algorithms (U.S. Food and Drug Administration, 2024). However, on April 2nd, 2019, the FDA published a paper that explains a potential pathway of how AI and ML technologies can obtain premarket review (U.S. Food and Drug Administration, 2024). My analysis can help guide this approval process and help create regulatory policies regarding the integration of ML and other AI technologies into healthcare. However, further analysis needs to be conducted in order to ensure that AI technologies are integrated into healthcare in a manner that protects privacy and upholds the ethical framework developed by the medical community.

One limitation of my analysis is that most of my research centered around three groups: software engineers, physicians, and patients. The perspectives of other relevant social groups such as the regulatory body also need to be researched. Another limitation of my study is that I did not discuss how ML healthcare applications impact a patient's privacy. Additionally, my analysis only focused on the four pillars of medical ethics. The principles of informed consent, truth-telling, and confidentiality all arise from the autonomy principle; however, I did not discuss these principles in my paper (Varkey, 2021). Despite these limitations, my analysis can still help individuals gain a better understanding of how ML algorithms can impact medical ethics.

References

Aksoy, S., & Tenik, A. (2002). The “four principles of bioethics” as found in 13th century Muslim scholar Mawlana’s teachings. *BMC Medical Ethics*, *3*(1), 4. https://doi.org/10.1186/1472-6939-3-4

AlHasan, A. (2023, February 8). Ethical Concerns Grow as AI Takes on Greater Decision-Making Role. *American College of Surgeons*. https://www.facs.org/for-medical-professionals/news-publications/news-and-articles/bulletin/2023/february-2023-volume-108-issue-2/ethical-concerns-grow-as-ai-takes-on-greater-decision-making-role/

Amisha, Malik, P., Pathania, M., & Rathaur, V. (2019). Overview of artificial intelligence in medicine. *Journal of Family Medicine and Primary Care*, *8*(7), 2328. https://doi.org/10.4103/jfmpc.jfmpc_440_19

Bakalova, Z., & Daneva, M. (2011). A comparative case study on clients participation in a “traditional” and in an Agile software company. *Proceedings of the 12th International Conference on Product Focused Software Development and Process Improvement*, 74–80. https://doi.org/10.1145/2181101.2181118

Benjamens, S., Dhunnoo, P., & Meskó, B. (2020). The state of artificial intelligence-based FDA-approved medical devices and algorithms: An online database. *Npj Digital Medicine*, *3*(1), 118. https://doi.org/10.1038/s41746-020-00324-0

Char, D. S., Abràmoff, M. D., & Feudtner, C. (2020). Identifying Ethical Considerations for Machine Learning Healthcare Applications. *The American Journal of Bioethics*, *20*(11), 7–17. https://doi.org/10.1080/15265161.2020.1819469

Char, D. S., Shah, N. H., & Magnus, D. (2018). Implementing Machine Learning in Health Care—Addressing Ethical Challenges. *New England Journal of Medicine*, *378*(11), 981–983. https://doi.org/10.1056/NEJMp1714229

Clifford, C. (2018, March 13). *Elon Musk: 'Mark my words — A.I. is far more dangerous than nukes.'* CNBC. https://www.cnbc.com/2018/03/13/elon-musk-at-sxsw-a-i-is-more-dangerous-than-nuclear-weapons.html

Dilsizian, S. E., & Siegel, E. L. (2014). Artificial Intelligence in Medicine and Cardiac Imaging: Harnessing Big Data and Advanced Computing to Provide Personalized Medical Diagnosis and Treatment. *Current Cardiology Reports*, *16*(1), 441. https://doi.org/10.1007/s11886-013-0441-8

Ding, Y., Sohn, J. H., Kawczynski, M. G., Trivedi, H., Harnish, R., Jenkins, N. W., Lituiev, D., Copeland, T. P., Aboian, M. S., Mari Aparici, C., Behr, S. C., Flavell, R. R., Huang, S.-Y., Zalocusky, K. A., Nardo, L., Seo, Y., Hawkins, R. A., Hernandez Pampaloni, M., Hadley, D., & Franc, B. L. (2019). A Deep Learning Model to Predict a Diagnosis of Alzheimer Disease by Using ^{18} F-FDG PET of the Brain. *Radiology*, *290*(2), 456–464. https://doi.org/10.1148/radiol.2018180958

Feudtner, C., Schall, T., Nathanson, P., & Berry, J. (2018). Ethical Framework for Risk Stratification and Mitigation Programs for Children With Medical Complexity. *Pediatrics*, *141*(Supplement_3), S250–S258. https://doi.org/10.1542/peds.2017-1284J

Gianfrancesco, M. A., Tamang, S., Yazdany, J., & Schmajuk, G. (2018). Potential Biases in Machine Learning Algorithms Using Electronic Health Record Data. *JAMA Internal Medicine*, *178*(11), 1544. https://doi.org/10.1001/jamainternmed.2018.3763

Gijsberts, C. M., Groenewegen, K. A., Hoefer, I. E., Eijkemans, M. J. C., Asselbergs, F. W., Anderson, T. J., Britton, A. R., Dekker, J. M., Engström, G., Evans, G. W., De Graaf, J., Grobbee, D. E., Hedblad, B., Holewijn, S., Ikeda, A., Kitagawa, K., Kitamura, A., De Kleijn, D. P. V., Lonn, E. M., … Den Ruijter, H. M. (2015). Race/Ethnic Differences in the Associations of the Framingham Risk Factors with Carotid IMT and Cardiovascular Events. *PLOS ONE*, *10*(7), e0132321. https://doi.org/10.1371/journal.pone.0132321

Graham, G. (2015). Disparities in Cardiovascular Disease Risk in the United States. *Current Cardiology Reviews*, *11*(3), 238–245. https://doi.org/10.2174/1573403X11666141122220003

Korkala, M., Abrahamsson, P., & Kyllonen, P. (2006). A Case Study on the Impact of Customer Communication on Defects in Agile Software Development. *AGILE 2006 (AGILE’06)*, 76–88. https://doi.org/10.1109/AGILE.2006.1

Ledford, H. (2019, October 24). *Millions of black people affected by racial bias in health-care algorithms*. Nature News. https://www.nature.com/articles/d41586-019-03228-6

Nadimpalli, M. (2017). *Artificial Intelligence Risks and Benefits*. *6*.

Nishida, N., Yamakawa, M., Shiina, T., Mekada, Y., Nishida, M., Sakamoto, N., Nishimura, T., Iijima, H., Hirai, T., Takahashi, K., Sato, M., Tateishi, R., Ogawa, M., Mori, H., Kitano, M., Toyoda, H., Ogawa, C., Kudo, M., & JSUM A. I. investigators. (2022). Artificial intelligence (AI) models for the ultrasonographic diagnosis of liver tumors and comparison of diagnostic accuracies between AI and human experts. *Journal of Gastroenterology*, *57*(4), 309–321. https://doi.org/10.1007/s00535-022-01849-9

Pinch, T. J., & Bijker, W. E. (1984). The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology might Benefit Each Other. *Social Studies of Science*, *14*(3), 399–441. https://doi.org/10.1177/030631284014003004

Rajkomar, A., Hardt, M., Howell, M. D., Corrado, G., & Chin, M. H. (2018). Ensuring Fairness in Machine Learning to Advance Health Equity. *Annals of Internal Medicine*, *169*(12), 866. https://doi.org/10.7326/M18-1990

Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., Lungren, M. P., & Ng, A. Y. (2017). *CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning* (arXiv:1711.05225). arXiv. http://arxiv.org/abs/1711.05225

Topff, L., Ranschaert, E. R., Bartels-Rutten, A., Negoita, A., Menezes, R., Beets-Tan, R. G. H., & Visser, J. J. (2023). Artificial Intelligence Tool for Detection and Worklist Prioritization Reduces Time to Diagnosis of Incidental Pulmonary Embolism at CT. *Radiology: Cardiothoracic Imaging*, *5*(2), e220163. https://doi.org/10.1148/ryct.220163

U.S. Food and Drug Administration. (2024, March 15). *Artificial Intelligence and Machine Learning in Software as a Medical Device*. U.S. Food and Drug Administration. https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device

Varkey, B. (2021). Principles of Clinical Ethics and Their Application to Practice. *Medical Principles and Practice: International Journal of the Kuwait University, Health Science Centre*, *30*(1), 17–28. https://doi.org/10.1159/000509119

Veatch, R. M. (1997). *Medical ethics* (2nd ed). Jones and Bartlett Publishers.