A Policy Analysis of Regulations on Artificial Intelligence for the Diagnosis of Disease

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

Samantha Clayton
Spring, 2020

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

A Policy Analysis of Regulations on Artificial Intelligence for the Diagnosis of Disease

Misdiagnosis: A Medical Crisis?

Most Americans will receive a medical misdiagnosis at least once in their lifetime (Rue, 2019). Medical misdiagnoses occur in 5% of outpatient office visits, 10% of hospital inpatient deaths, and 12% of hospital adverse events, and contribute to 74,000 deaths per year in the United States (U.S.) (Papier, 2015). Artificial intelligence (AI) will decrease the rate of misdiagnoses. AI is defined as "a broad discipline that aims to understand and design systems that display properties of intelligence, emblematic of which is the ability to learn, or derive knowledge from data" (Panch et al., 2018). Significant research over the past few decades aims at using AI for medical and healthcare applications, including surgery, administration, and as nursing assistants. AI will be utilized to help with the diagnosis of diseases by revealing previously hidden trends in data, and thus will have substantial impact both at the individual patient and system level (Panch et al., 2018).

This technology is still in the developmental stage, but has prompted concerns from health policy scholars as to the possible risks of its implementation. These systems prompt concerns related to the privacy of patient data, the quality and safety of the algorithms, and their impacts on the role of physicians and the healthcare system at large. Due to the threat these risks pose, AI for the diagnosis of diseases must be considered as an inherently political technology. This classification is in part due to how its development may conflict with current regulations and the need for government and healthcare systems to protect the safety and privacy of patients. The goal of this thesis to identify the need for the regulation of medical diagnoses AI to protect patient privacy and promote patient care.

Research Methods

The question that is the focus of this research paper is: where is the need for future policies surrounding AI used for diagnosis to protect patient privacy while providing the best possible care. This research utilizes documentary research methods and policy analysis to consider the perceptions of this technology in the scientific community, medical community, political space, and general public. Documentary research methods are implemented by an extensive literature review to understand the current state of this technology and the successes it has had, the potential this technology has to grow into in the future, and the risks produced by its adoption into healthcare. This research utilizes databases, such as PubMed, Web of Science, Congress.gov, and Public Affairs Information Service International (PAIS Index). In addition, journals, such as the AMA Journal of Ethics, Health Affairs, and Nature, were used to understand the perspectives of the medical and scientific community, along with studying proposed frameworks for the development of these systems. Policy analysis was used to analyze the regulations that have been applied to the use of AI as a whole, and specifically in the healthcare system. This analysis looked at current policies, such as Health Insurance Portability and Accountability Act (HIPAA) of 1996 and the American Medical Association's (AMA's) policy on Augmented Intelligence released in 2018 (Crigger & Khoury, 2019; Office for Civil Rights, 2009). Other policies, such as the proposed Algorithmic Accountability Act of 2019 and the California Consumer Privacy Act, were used to understand how AI generally is being addressed through policy, and how these policies could be fit to AI for diagnosis. The analysis of this legislation focuses separately on the regulations aimed at improving each of the limitations or risks of AI.

The Rise of AI in Healthcare

Approximately 12,000 U.S. patients experience a misdiagnosis per year (Singh et al., 2014). While the use of AI poses obvious risks to a person's health, there is also a substantial impact on the healthcare system at large. However, AI systems can leverage the vast amount of health data to produce accurate diagnoses for a variety of conditions and potentially reduce the risk of misdiagnoses. The great potential of AI to improve the quality of medical diagnoses has been a draw for venture funding with 19 private digital health companies receiving $330.4 million from 2011-2017 (Zweig & Tran, 2018). In 2017, analysts predicted that AI applied for preliminary and automated image diagnosis could save $8 billion by 2026 from healthcare spending (HealthITAnalytics, 2019). While AI has yet to be implemented in hospitals with the purpose of diagnosis, the research has been promising. In a review of studies, AI programs performed as well as, or slightly better, than human healthcare professionals overall (Liu et al., 2019). Just this year, Google's Deep Mind AI outperformed radiologists in diagnosing breast cancer by 11% (McKinney et al., 2020)While there is a need for further improvement in the accuracy of these systems before widespread use, these systems have already vastly improved from their initial development during the 1980s. The need for further development leaves time for regulations and policy to be implemented to address the future of AI for diagnosis.

While AI systems can improve diagnostic accuracy, they have several weaknesses that need to be addressed. For example, the strength of a given AI-drive algorithm depends on the data with which it works. Developers must carefully choose the data sets for these training systems. Otherwise AI systems can, invisibly and unintentionally, reproduce or magnify the biases of the engineers designing them that may risk exacerbating existing racial health disparities (Crigger & Khoury, 2019). When AI has been applied to other disciplines, a racial bias has been observed in the results (Obermeyer et al., 2019). In addition, the requirement of

such large amounts of data for effective predictions, will mean patients will need to open up their data to wider use, thus relinquishing some of their privacy (Longhurst et al., 2014). Finally, these systems will represent a change to the duties of physicians. Whether they are used only as a supplementary tool or their predictions are given more weight, physicians will need to be trained differently to prepare for this change (Cabitza et al., 2017). Scientific policy offers a route to mitigate these potential dangers associated with the use of AI in time for the implementation of this technology.

Is Diagnostic AI Political?

The field of STS has examined the use of AI from many directions. This thesis mainly focuses on examining the public effect of the technology, especially in regards to how policy should address these effects. To that end, STS scholars published in the journal, Science and Public Policy, have addressed concerns over the impacts of AI generally. For instance, Fred Phillips, a current professor at University of New Mexico and expert on technology management research, notes that AI represents an advancement in the methodology for solving problem, which in turn drives advances in scientific progress (Phillips, 2017). However, he goes further to relay the importance that scientists "ensure that big data's benefits outweigh any possible harm, and that we build on big data's strengths while shoring up its weaknesses" (Phillips, 2017). Meanwhile, Marie Jahoda, an Austrian-British social psychologist, had a more skeptical view of the future of AI, encouraging the use of AI as just a tool and never the final decision-maker (Jahoda, 1986). Yet, even with her concerns over the dangers of AI, she believed that the research into AI should not be restricted to a clearly known set of goals (Jahoda, 1986). An open approach to the use of AI in research led to its use in a vast number of fields today, presenting further benefits to its continued development.

This thesis utilizes the STS framework of political technologies to examine the use of AI systems to diagnose diseases. The theory of political technologies considers the political nature of technologies. Artifacts can be considered political either because they were developed with the intention of having political effects or the device itself has unintentional political impacts. The latter group is referred to as inherently political technologies. AI systems designed to diagnose disease fall into the latter category. While AI seeks to solve medical problems, the implementation in health care systems would require the consideration of a patient's right to privacy, unintended impacts on underrepresented races or classes, and a potential shift to the role of a physician within the healthcare system. The use of the framework of political technologies presents some risk of critique. While this theory can be useful for examining the interactions between a new technology and the political system, it has been criticized in the past for being applied too generously, to technologies that are actually apolitical or inaccurately judged (Joerges, 1999). However, as will be shown in this thesis, the impacts of this technology will affect the healthcare landscape in such a way that this view is necessary.

Do Current Policies Protect Us?

The U.S. has extensive regulations that pertain to the safety and care for electronic health data to maintain the privacy of patients. However, AI presents a more extensive need for data that challenges these current regulations. The current legislation was developed before there was a major focus on using AI in healthcare beyond research purposes. To allow for the use of patient data on a widespread basis for the development and use of AI systems, there will need to be an adjustment to existing laws. Yet, AI also presents risk to minority populations if legislation or policies are not implemented to control for bias in these AI systems. Finally, AI for diagnosis prompts some confusion as to who is liable when mistakes occur that needs to be laid out.

One of the main concerns that the use of AI for the diagnosis of disease prompts is whether this use conflicts with current laws designed to protect patient privacy. The main law in the U.S. that addresses the privacy of health data is HIPAA. This act establishes national standards to protect individuals' medical records and other personal health information and sets limits and conditions on the uses and disclosures that may be made of such information without patient authorization (Office for Civil Rights, 2009). This act creates two pathways for how patient data is allowed to be used in research. Individually identifiable health information can be used for research purposes with patient consent. Meanwhile, there are no restrictions on the use or disclosure of de-identified health information, which does not identify nor provide reasonable basis to identify an individual. However, current technology blurs the line between these two data types. The ability to identify an individual without the usual identifiers has become easier. (Crigger & Khoury, 2019). Due to this inability to have truly unidentifiable data, the requirements to gain consent for the use of medical data should be expanded further. By creating conflict with current laws and policy paradigms, diagnostic AI reaffirms its classification as a political technology. Further efforts will require steps both at the state and national level that aim to expand on the need for consent into the private sector.

The proposed California Consumer Privacy Act would grant a consumer the right to request a business to disclose the categories and specific pieces of personal information that it collects about the consumer, the categories of sources from which that information is collected, the business purposes for collecting or selling the information, and the categories of third parties with which the information is shared (Palsan, 2018). The law would allow consumers to know if companies developing AI systems are using their data and give them the opportunity to opt out of its use. In addition, the Fundamentally Understanding The Usability and Realistic Evolution

(FUTURE) of Artificial Intelligence Act of 2017 and the Mind Your Own Business Act of 2019 have been introduced in Congress that will further examine the impact of AI on the right to privacy and expand the expectations for privacy protection (Delaney, 2018; Wyden, 2019).

Countries such as Brazil and Canada may serve as good models for which the United States can learn from. Brazil and Canada's recent expansion of privacy laws increasing the stringency of the requirement for consent for the use of data. (Ahmad, 2019). A recent United Kingdom survey reports that 63% of the adult population is uncomfortable with allowing personal data to be used to improve healthcare and is unfavorable to AI systems replacing doctors and nurses in tasks they usually perform (Vayena et al., 2018). To allow patients to have more control over how their data is used that is in line with these views, the European Union (EU) has an expansive privacy law, known as the General Data Protection Regulation (GDPR), which requires all data protection and use to be opt-in, including requiring explicit consent for research use of health data only where research serves a "high public interest" (Pesapane et al., 2018). While the EU law makes great strides in giving patients control over their data, the GDPR may go too far, slowing down research or complicating the ability of researchers to gather substantial data. AI is inherently political as its use demands intervention by the government to decide how and when it can use patient data. Therefore, the U.S. should look to increase the need for consent for research without going as far as to decide what type of research is acceptable.

Beyond regulating how this data is used, developers need to be certain that they keep this data safe from hackers. The expectations for the security of health data systems was established with HIPAA and expanded with the Health Information Technology for Economic and Clinical Health (HITECH) Act, Section 5 of the Federal Trade Commission (FTC) Act, and the 21st Century CURES Act (Bonamici, 2016; Ohlhausen, 2014; Rights (OCR), 2009). These laws

broadened the data protection requirements of HIPAA, gave the FTC the authority to enforce present regulations, and defined security requirements for the transfer of data between institutions. In addition, the proposed Securing American Leadership in Science and Technology Act of 2020 would support the development of new cybersecurity and data encryption tools that could be used to protect data (Lucas, 2020). Another technical solution that has been proposed by the American Medical Association (AMA) is the use of block-chain style technologies to secure data and track access (Crigger & Khoury, 2019). The U.S. has well-established regulations for data protection that are flexible enough to adapt to the implementation of AI systems in healthcare, but should ensure that these data protection systems keep up with the advancement in adversarial technologies.

Another major concern as the use of AI expands to more fields, and especially as it will participate in health decisions, such as in the case of its use for diagnosis, is its potential to blindly perpetuate bias. There are several proposed bills in Congress addressing this problem, such as the Algorithmic Accountability Act of 2019, the FUTURE of Artificial Intelligence Act of 2017, and the Growing Artificial Intelligence Through Research (GrAITR) Act (Clarke, 2019; Delaney, 2018; Lipinski, 2019). While the latter two proposed bills aim to study the potential for bias in these systems and advise careful development of AI to avoid bias, the Algorithmic Accountability Act takes the strongest steps to combat bias by requiring each entity to produce an automated decision system impact assessment that would evaluate the design and training data for impacts on accuracy, fairness, bias, discrimination, privacy, and security (Clarke, 2019). The Algorithmic Accountability Act places the responsibility firmly on the developers of these systems to assess its potential impact, yet is restricted to large companies and ambiguous about what these assessments would look like in the iterative process of software development (New,

2019). Further explanation of this process would make the process clearer for companies looking to develop this technology in a just way. In addition, connecting this legislation with liability, could ensure that companies complete these assessments thoroughly allowing for the most fair and effective technology to be created. Finally, reports from multiple groups, including the AMA, the National Institute for Standards and Technology (NIST), and the National Science and Technology Council have supported the idea of creating open repositories and promoting data sharing, so that the datasets used in developing AI systems can be grown and examined for any issues or biases (Crigger & Khoury, 2019; National Institute for Science and Technology, 2019; National Science and Technology Council, 2016). The implementation of a revised Algorithmic Accountability Act and the promotion of data sharing policies will take major steps towards mitigating the risk of bias in AI used for the diagnosis of disease.

Limitations & Future Research

There are a number of proposed bills and newly minted committees that aim to limit the harm that could be introduced by widespread implementation of AI in a variety of fields. However, much of the focus is on what private companies are doing with our data and could do in the future, thus, legislators have placed an emphasis on creating general legislation to impact the private sector. Yet, one of the areas that could benefit the most from the implementation of AI is the healthcare industry. Laws designed for the private sector cannot necessarily be applied directly to AI for diagnosis due to the sensitivity of the medical data and the complication of the existing regulations and bureaucracy of the healthcare system.

In addition to the unique framework of AI in healthcare, the current rapid speed development and new implementation of AI has meant that laws aimed at addressing this technology are relatively new or not yet implemented. Thus, it is too early to judge the impacts

these laws could have on the development of the industry. Artificial intelligence for the diagnosis of disease is in an especially early stage of development, thus the policies addressing it are still starting to take effect.

After the legislation has been put into effect, it will be important that researchers are examining the impacts of the new legislation. The researchers must also look for additional threats that emerge as AI becomes a larger part of our lives. Beyond just understanding the effects of these policies on the safety and privacy of patients, further research should examine the impact of these policy changes on the perceptions of AI for the diagnosis of disease held by the public, members of the scientific community, and members of the healthcare system. It is incumbent on researchers and industry leaders to suggest the best path forward to allow this technology to fulfill its potential to benefit the lives of patients.

Conclusion

Alterations to the current laws surrounding privacy and data protection and the advancement of current proposals to mitigate the risk of bias in AI would lower the harm that the implementation of AI for diagnosis could have on patients. These changes to regulations would allow AI to be implemented with minimal risk, and thus decrease the rate of misdiagnoses to save lives. While it will be several years before the impact of these laws are known, the rapid changing landscape of this field requires that proactive action is taken to mitigate the harm of AI for the diagnosis of disease. By accounting for the risks of this technology through policy changes, patients will be able to benefit from this technology without giving up their right to privacy or causing them further harm.

References

Ahmad, T. S. (2019, January). *Regulation of artificial intelligence* [Web page]. https://www.loc.gov/law/help/artificial-intelligence/americas.php

Bonamici, S. (2016, December 13). *H.R.34 - 114th Congress (2015-2016): 21st Century Cures Act* [Webpage]. https://www.congress.gov/bill/114th-congress/house-bill/34

Cabitza, F., Rasoini, R., & Gensini, G. F. (2017). Unintended consequences of machine learning in medicine. *JAMA*, *318*(6), 517–518. https://doi.org/10.1001/jama.2017.7797

Clarke, Y. D. (2019, April 10). *H.R.2231 - 116th Congress (2019-2020): Algorithmic Accountability Act of 2019* [Webpage]. https://www.congress.gov/bill/116th-congress/house-bill/2231

Crigger, E., & Khoury, C. (2019). Making policy on augmented intelligence in health care. *AMA Journal of Ethics*, *21*(2), 188–191. https://doi.org/10.1001/amajethics.2019.188.

Delaney, J. K. (2018, May 22). *Text - H.R.4625 - 115th Congress (2017-2018): FUTURE of Artificial Intelligence Act of 2017* [Webpage]. https://www.congress.gov/bill/115th-congress/house-bill/4625/text

HealthITAnalytics. (2019, July 29). *Future AI opportunities for improving care delivery, cost, and efficacy*. HealthITAnalytics. https://healthitanalytics.com/news/future-ai-opportunities-for-improving-care-delivery-cost-and-efficacy

Jahoda, M. (1986). Artificial intelligence: An outsider's perspective. *Science and Public Policy*, *13*(6), 333–340. https://doi.org/10.1093/spp/13.6.333

Joerges, B. (1999). Do Politics Have Artefacts? *Social Studies of Science*, *29*(3), 411–431. https://doi.org/10.1177/030631299029003004

Lipinski, D. (2019, April 10). *Text - H.R.2202 - 116th Congress (2019-2020): GrAITR Act* [Webpage]. https://www.congress.gov/bill/116th-congress/house-bill/2202/text

Liu, X., Faes, L., Kale, A. U., Wagner, S. K., Fu, D. J., Bruynseels, A., Mahendiran, T., Moraes, G., Shamdas, M., Kern, C., Ledsam, J. R., Schmid, M. K., Balaskas, K., Topol, E. J., Bachmann, L. M., Keane, P. A., & Denniston, A. K. (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: A systematic review and meta-analysis. *The Lancet Digital Health*, *1*(6), e271–e297. https://doi.org/10.1016/S2589-7500(19)30123-2

Longhurst, C. A., Harrington, R. A., & Shah, N. H. (2014). A 'green button' for using aggregate patient data at the point of care. *Health Affairs*, *33*(7), 1229–1235. https://doi.org/10.1377/hlthaff.2014.0099

Lucas, F. D. (2020, January 28). *Text - H.R.5685 - 116th Congress (2019-2020): Securing American Leadership in Science and Technology Act of 2020* [Webpage]. https://www.congress.gov/bill/116th-congress/house-bill/5685/text

McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., Back, T., Chesus, M., Corrado, G. C., Darzi, A., Etemadi, M., Garcia-Vicente, F., Gilbert, F. J., Halling-Brown, M., Hassabis, D., Jansen, S., Karthikesalingam, A., Kelly, C. J., King, D., … Shetty, S. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, *577*(7788), 89–94. https://doi.org/10.1038/s41586-019-1799-6

National Institute for Science and Technology. (2019). *U.S. leadership in AI: A plan for federal engagement in developing technical standards and related tools* (p. 52). https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf

National Science and Technology Council. (2016, May 3). *Preparing for the future of artificial intelligence*. Whitehouse.Gov. https://obamawhitehouse.archives.gov/blog/2016/05/03/preparing-future-artificial-intelligence

New, J. (2019, September 23). *How to fix the Algorithmic Accountability Act*. Center for Data Innovation. https://www.datainnovation.org/2019/09/how-to-fix-the-algorithmic-accountability-act/

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, *366*(6464), 447–453. https://doi.org/10.1126/science.aax2342

Office for Civil Rights. (2009, November 20). *Summary of the HIPAA Security Rule* [Text]. HHS.Gov. https://www.hhs.gov/hipaa/for-professionals/security/laws-regulations/index.html

Ohlhausen, M. K. (2014). Section 5 of the FTC Act: Principles of navigation. *Journal of Antitrust Enforcement*, *2*(1), 1–24. https://doi.org/10.1093/jaenfo/jnt013

Palsan, S. (2018). The California Consumer Privacy Act: Getting a head start on compliance. *Ogletree Deakins*. https://ogletree.com/insights/2018-10-08/the-california-consumer-privacy-act-getting-a-head-start-on-compliance

Panch, T., Szolovits, P., & Atun, R. (2018). Artificial intelligence, machine learning and health systems. *Journal of Global Health*, *8*(2). https://doi.org/10.7189/jogh.08.020303

Papier, A. (2015). *Improving diagnostic accuracy*. VisualDX. https://www.visualdx.com/learn-more/improving-diagnostic-accuracy

Pesapane, F., Volonté, C., Codari, M., & Sardanelli, F. (2018). Artificial intelligence as a medical device in radiology: Ethical and regulatory issues in Europe and the United States. *Insights into Imaging*, *9*(5), 745–753. https://doi.org/10.1007/s13244-018-0645-y

Phillips, F. (2017). A perspective on 'Big Data.' *Science and Public Policy*, *44*(5), 730–737. https://doi.org/10.1093/scipol/scx012

Rights (OCR), O. for C. (2009, October 28). *HITECH Act Enforcement Interim Final Rule* [Text]. HHS.Gov. https://www.hhs.gov/hipaa/for-professionals/special-topics/HITECH-act-enforcement-interim-final-rule/index.html

Rue, N. (2019, January 30). *AI and healthcare: The battle against misdiagnosis*. Medium. https://becominghuman.ai/ai-and-healthcare-the-battle-against-misdiagnosis-7eb3742ad1f7

Singh, H., Meyer, A. N. D., & Thomas, E. J. (2014). The frequency of diagnostic errors in outpatient care: Estimations from three large observational studies involving US adult populations. *BMJ Quality & Safety*, *23*(9), 727–731. https://doi.org/10.1136/bmjqs-2013-002627

Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLoS Medicine*, *15*(11). https://doi.org/10.1371/journal.pmed.1002689

Wyden, R. (2019, October 17). *Text - S.2637 - 116th Congress (2019-2020): Mind Your Own Business Act of 2019* [Webpage]. https://www.congress.gov/bill/116th-congress/senate-bill/2637/text

Zweig, M., & Tran, D. (2018). *The AI/ML use cases investors are betting on in healthcare*. Rock Health. https://rockhealth.com/reports/the-ai-ml-use-cases-investors-are-betting-on-in-healthcare/

Zweig, M., & Tran, D. (2018). *The AI/ML use cases investors are betting on in healthcare*. Rock Health. https://rockhealth.com/reports/the-ai-ml-use-cases-investors-are-betting-on-in-healthcare/