Solving Pay Stub Compliance Issues Through Asynchronous Generation and Persisting to S3
(Technical Topic)

Investigating the Effects of Algorithmic Bias in Lung Cancer Diagnosis and Treatment
(STS Topic)

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Preethi Chidambaram

December 2, 2022

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.

ADVISORS

MC Forelle, Department of Engineering and Society

Rosanne Vrugtman and Briana Morrison, Computer Science

Introduction

One of the main uses for artificial intelligence involves processing large amounts of data and smart decision making based on patterns within datasets (Brooks, 2019). In the healthcare industry, AI is used to manage millions of patient records and analyze data to select customized treatment plans for patients (Basu et al., 2020). Despite the benefits it brings in increasing hospital efficiency, AI poses considerable risks to patient care, privacy, and social bias. There is a common misconception surrounding bias in AI because people automatically assume that human based decisions are inherently more biased than computer-based decisions due to the factor of emotion. However, it should be noted that AI systems depend heavily on the data available to them, and thus can unintentionally perpetuate social biases already present in society (Manyika et al., 2019). As an engineer, I believe it is important to explore the social consequences of data management within automated systems because it can significantly affect the decisions that users of these automated systems make.

The performance of artificially intelligent software heavily depends on the data used to train the models. In the case a particular disease has little to no data, AI models cannot be trusted to make accurate decisions regarding patient health (*Pros & Cons of Artificial Intelligence in Medicine*, 2021). AI models have the tendency to overlook social variables when coming to a medical decision (McKendrick & Thurai, 2022). While AI systems may be able to accurately diagnose diseases, they may not account for economic restrictions or personalized preferences of a patient when creating a treatment plan for them (*Pros & Cons of Artificial Intelligence in Medicine*, 2021). Research on algorithmic bias in the healthcare industry warrants more significance because it will help identify sources of implicit bias that contribute to racial and socioeconomic discrimination. My STS topic will delve into the effects of algorithmic bias in the

treatment of lung cancer amongst underrepresented social groups. By pursuing this research topic, I intend to find flaws in the development of medical AI training datasets and help establish a set of standards that must be followed in the development of these systems.

The technical research portion of this thesis addresses the social issues that arise when data is not properly recorded and maintained. It is motivated by the fact that a company can face problems concerning liability when their system does not have adequate measures in place for tracking the version history of documents (Kolodziej, 2021). Specifically, my technical project aims to solve compliance issues within Gusto's payroll system by creating a backend data model to store pay stub file version history. As a payroll and benefits management platform, Gusto's product deals with the movement of valuable commodities, which means they must prioritize transparency of data records to their customers. My technical and STS topics both examine the ethical consequences of technical systems that are built on inaccurate data models.

Technical Topic

Gusto is a company that provides cloud based services to small business for managing payroll, employee benefits, and human resources (*Welcome to the Gusto People Platform*, n.d.). Automated business management systems have shaped the way companies are run by making the logistical tasks of onboarding and paying employees more efficient and less prone to human error (Zaric, 2021). Mahajan et al. (2015) provide an overview of the components contained in a basic digital payroll system, which helped inform my technical work.

Gusto's payroll management platform currently does not keep track of pay stub version history because pay stubs are generated on demand, meaning that they are only generated upon requests from the user. The pay stub file is created by aggregating the most up to date information stored in the payroll data table from the associated pay period. This is acceptable for

most cases; however, it gets tricky if any data is modified after the original payroll was processed. For example, when a reversal payroll is run, the system generates a pay stub which may not be an accurate representation of what the employee was actually paid at the time of running the payroll (*Gusto Help Center - Edit, Cancel, or Reverse an Employee Payroll or Contractor Payment*, n.d.). Another possible case that needs attention is when an employee logs paid time off (PTO) after the payroll is processed and the system assumes it should deduct the amount of PTO from the employee's regular hours worked. The pay stub file will be automatically updated to reflect the changes made to the payroll, but it overwrites the preexisting pay stub file without storing it anywhere.

Gusto is considered to be non-compliant due to its failure to maintain an accurate paper trail of employee pay when a payroll was originally run. If a customer were to specifically request a pay stub from the time of running the payroll, Gusto is required to provide it. Since the technical system cannot support this behavior, Gusto relies on customer support lines to retrieve the correct information and generate the pay stub file manually for the customer.

To address this compliance issue, I created a backend model to store pay stubs and persist them to S3, Amazon's storage service. I structured the design of this model similar to existing models containing pdf attachments. I also modified the on demand pdf generation behavior to be asynchronous to improve overall performance. I decided to store pay stub pdf files in S3 because it seemed to be a simple, cost-efficient solution to tracking version history of documents (Midgeley, 2022). While restructuring the payroll database, I found that preserving aspects of the existing pay stub behavior was important for maintaining a satisfactory customer experience. While the backend work for this project was completed, the frontend system for displaying version history to the user is still being developed.

STS Topic

My STS research, in contrast with my technical research, will investigate the effects of poor design decisions in AI models on the treatment and care of patients diagnosed with lung cancer. Artificial intelligence has improved several aspects of the healthcare industry by streamlining procedures and efficiently analyzing patient records. Specifically with regards to lung cancer diagnosis, AI recognition technology can accurately process images to help doctors screen for early stages of cancer (Pei et al., 2022). AI systems are also used in precision medicine to build personalized treatment plans for patients by automatically extracting features from medical data (Chen et al., 2021).

Despite the advancements it brings, there are many factors that cause AI models to make unfair decisions and perpetuate social biases revolving around race and socioeconomic status, which has potential legal implications. Under the Equality Act 2010, "discriminatory decisions that are mediated by an apparently neutral provision, criterion or practice, but still disproportionately affect certain protected groups are qualified as indirect discriminations" (Schönberger, 2019, 184). This anti-discrimination law applies to medical AI systems because these systems are designed to be neutral entities that make decisions regarding public health (Schönberger, 2019). In order to decrease the chances of discrimination within AI systems, developers must first agree upon a universal definition for terms like fairness and bias and then translate them into computational parameters. Paulus & Kent (2020) discuss that fairness in AI refers to treating individuals equally despite certain protected characteristics. On the other hand, bias in AI refers to problems with the model design that may disproportionately affect performance in relation to certain groups of people (Paulus & Kent, 2020). The main contributing factor of bias is when a model's training data is racially skewed due to the

underrepresentation of certain groups of people (Dankwa-Mullan & Weeraratne, 2022). For example, an AI algorithm used to identify patients for extra care based its decisions on a person's medical costs even though the relationship between healthcare needs and cost can vary greatly (Obermeyer et al., 2019). This design factor resulted in the algorithm making decisions where black patients needed to reach a higher threshold than white patients to be considered for extra care. Obermeyer et al. (2019) concluded that the data labels were the key determinants of this AI algorithm's predictive quality and bias and they showed that redesigning the dataset can change the results. This study and several others highlight the need to establish a set of standards regarding how AI training datasets are built to ensure that diversity and inclusion are prioritized.

I will use the STS framework Actor Network Theory to analyze the relationship between patients, doctors, and the AI algorithms used to predict diagnosis and treatment plans for patients. Actor network theory (ANT) claims human and non-human actors are enrolled in the construction of technological systems while also emphasizing generalized symmetry, meaning these human and non-human actors must be treated equally (Cressman, 2009 as cited in Latour, 2005). I intend to apply this theory to analyze how work is delegated to AI systems to better understand the role of AI in oncology. Past studies have shown that humans collaborating with AI systems to formulate patient diagnoses and treatment plans result in more inclusive decisions (Banks, 2018). Taking this into account, it is important to identify developmental aspects of these systems that might restrain the way humans interact with them (Woolgar, 1990). I hope to use ANT to further study this relationship between medical professionals and patients, the human actors, and the AI systems, the non-human actor.

Research Questions and Methods

My proposed research question is: *How has algorithmic bias in lung cancer diagnosis and treatment negatively impacted patients from underrepresented ethnic and socioeconomic groups?* I intend to use the research method of literature review by compiling existing pieces of literature that focus on the development of AI systems within the field of cancer treatment. I have already found papers that detail the functions of AI in diagnosing patients with cancer, and I hope to narrow my literature search by looking at the applications of AI diagnosis systems in lung cancer. In addition to this thread, I will gather literature pieces that delve into the relationship between existing AI systems and patients from lower income or underrepresented ethnic groups. Searching for literature within this second thread will help me compile information about the social aspects of AI systems to further assess their effects on patients. By reviewing literature in these two threads, I hope to gain more insight into how the design of AI systems can impact the field of oncology.

Conclusion

My technical project aims to improve the persistence of data to ensure that users have access to reliable information. Maintaining accessibility and transparency of data records will help users be better informed when making decisions that affect their business. My technical work produced a backend feature improvement to Gusto's payroll database that keeps track of document version history. Additionally, through my STS research, I will analyze the development and effects of such AI systems to gain a better understanding of their relationship in the healthcare network. Using resources available to me, I will examine whether AI algorithms effectively integrate social aspects of patient history when making predictions about diagnoses and creating personalized treatment plans. By synthesizing information from various sources, I

hope to find ways in which AI systems can be improved by guiding the establishment of standards surrounding the development of these systems.

References

Banks, J. (2018). The Human Touch: Practical and Ethical Implications of Putting AI and Robotics to Work for Patients. *IEEE Pulse*, *9*(3), 15–18. https://doi.org/10.1109/MPUL.2018.2814238

Basu, K., Sinha, R., Ong, A., & Basu, T. (2020). Artificial Intelligence: How is It Changing Medical Sciences and Its Future? *Indian Journal of Dermatology*, *65*(5), 365–370. https://doi.org/10.4103/ijd.IJD_421_20

Brooks, A. (2019, November 4). *The Benefits of AI: 6 Societal Advantages of Automation*. https://www.rasmussen.edu/degrees/technology/blog/benefits-of-ai/

Chen, Z.-H., Lin, L., Wu, C.-F., Li, C.-F., Xu, R.-H., & Sun, Y. (2021). Artificial intelligence for assisting cancer diagnosis and treatment in the era of precision medicine. *Cancer Communications*, *41*(11), 1100–1115. https://doi.org/10.1002/cac2.12215

Cressman, D. (2009). *A Brief Overview of Actor-Network Theory: Punctualization, Heterogeneous Engineering & Translation*. Simon Fraser University. https://summit.sfu.ca/item/13593

Dankwa-Mullan, I., & Weeraratne, D. (2022). Artificial Intelligence and Machine Learning Technologies in Cancer Care: Addressing Disparities, Bias, and Data Diversity. *Cancer Discovery*, *12*(6), 1423–1427. https://doi.org/10.1158/2159-8290.CD-22-0373

Gusto Help Center—Edit, cancel, or reverse an employee payroll or contractor payment. (n.d.). Retrieved October 27, 2022, from https://support.gusto.com/article/106621949100000/Edit-cancel-or-reverse-an-employee-payroll-or-contractor-payment

Kolodziej, M. (2021, September 2). *The Five Risks of Poor Records Management*. Access. https://www.accesscorp.com/blog/the-five-risks-of-poor-records-management/

Latour, Bruno (2005). Reassembling the Social: An Introduction to Actor-Network Theory. New York: Oxford University Press.

Mahajan, K., Shukla, S., & Soni, N. (2015). A Review of Computerized Payroll System. *IJARCCE*, 67–70. https://doi.org/10.17148/IJARCCE.2015.4113

Manyika, J., Silberg, J., & Presten, B. (2019, October 25). What Do We Do About the Biases in AI? *Harvard Business Review*. https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai

McKendrick, J., & Thurai, A. (2022, September 15). AI Isn't Ready to Make Unsupervised Decisions. *Harvard Business Review*. https://hbr.org/2022/09/ai-isnt-ready-to-make-unsupervised-decisions

Midgeley, A. (2022, February 25). *The Essential Guide to AWS S3 Pricing—Apptio*. https://www.apptio.com/blog/essential-guide-aws-s3-pricing/

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

Paulus, J. K., & Kent, D. M. (2020). Predictably unequal: Understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. *Npj Digital Medicine*, *3*(1), 99. https://doi.org/10.1038/s41746-020-0304-9

Pei, Q., Luo, Y., Chen, Y., Li, J., Xie, D., & Ye, T. (2022). Artificial intelligence in clinical applications for lung cancer: Diagnosis, treatment and prognosis. *Clinical Chemistry and Laboratory Medicine*. https://doi.org/10.1515/cclm-2022-0291

Pros & Cons of Artificial Intelligence in Medicine. (2021, August 17). College of Computing & Informatics. https://drexel.edu/cci/stories/artificial-intelligence-in-medicine-pros-and-cons/

Schönberger, D. (2019). Artificial intelligence in healthcare: A critical analysis of the legal and ethical implications. *International Journal of Law & Information Technology*, *27*(2), 171–203. https://doi.org/10.1093/ijlit/eaz004

Welcome to the Gusto People Platform. (n.d.). Gusto. Retrieved October 27, 2022, from https://gusto.com

Woolgar, S. (1990). Configuring the user: The case of usability trials. *Sociological Review*, *38*(1), 58–99. https://doi.org/10.1111/j.1467-954X.1990.tb03349.x

Zaric, S. (2021, September 28). *Payroll Automation Guide: Definition, Advantages, and FAQs*. https://www.deel.com/blog/automated-payroll-system-advantages