

Corporate Adoption of AI Technology

The Use of Artificial Intelligence in Hiring and Recruitment

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

Liam Tolbert

November 3, 2023

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

Prof. Pedro Augusto P. Francisco, Department of Engineering and Society

Prof. Briana Morrison, Department of Computer Science

Introduction

AI is taking over the world. Companies everywhere are scrambling to adopt it as fast as possible and use it in as many places and ways as possible. Here, I detail my experience with this trend as well as a thorn in the side of companies preparing to utilize it to their fullest potential.

In my Capstone/Technical Project, I delved into the integration of artificial intelligence within a major government contractor's operations. During my internship, I was involved in projects that aimed to harness AI's potential for enhancing efficiency and security within the organization. One notable initiative centered on developing a Chrome extension to oversee employee interactions with generative AI platforms, ensuring the prevention of sensitive data exposure. Another project involved creating an in-house generative AI system, customized to serve as an internal knowledge base for employees. These projects were the company's attempt to take advantage of AI as a new tool to its furthest extent.

My STS research addressed the challenges surrounding fairness in AI/ML hiring and recruitment algorithms. These algorithms, increasingly adopted by companies to streamline candidate screening processes, promise efficiency and impartiality. However, my research exposed the inherent biases in these systems, arising from training data and algorithmic complexities. These biases have far-reaching consequences, affecting not only individual job applicants but also society at large by impacting workforce diversity and stifling innovation. The common thread binding these topics is the transformative power of AI, underpinning both corporate operations and societal equity. As the ever-accelerating wave of AI innovation sweeps through industries, it is impossible to overstate the significance of this technology. Its potential is seemingly limitless, and companies are racing to harness it in as many places and ways as

possible. From automating routine tasks to enabling breakthroughs in healthcare, education, and beyond, the transformative possibilities appear to have no bounds.

However, amidst the excitement of these possibilities lies a pressing need for responsibility and ethical consideration. The rapid adoption of AI, while promising remarkable improvements in efficiency, has also unearthed complex socio-technical challenges. These challenges encompass issues of bias, discrimination, and the socio-economic impact of AI's omnipresence, which affect not only individual job applicants, but also society at large. The core idea that emerges from this connection is the dual role AI plays in both enhancing corporate landscapes and influencing societal fairness. It underscores the ethical responsibilities of engineers and computer scientists as they navigate the uncharted territories of technology and its impact on society. As AI continues to redefine how we work, live, and interact, our collective responsibility to navigate these uncharted waters with ethical diligence becomes increasingly imperative.

Corporate Adoption of AI Technology

Over the summer, I worked as a software engineer intern for a large government contractor based in Northern Virginia. During the period where I worked full-time (as I am still employed by them part-time), it was probably some of the most I have ever learned about any topic I was passionate about in such a short period of time. In my internship, I was able to directly contribute to the (all unclassified) projects the senior developers were doing while being treated as a pseudo-junior developer by my manager in terms of guidance and task assignment.

When I started my employment, the organization was in its initial phases of exploring the integration of artificial intelligence into its daily operations. I had the privilege of participating in two significant projects during this period. The first project centered on the development of a

Chrome extension aimed at supervising and regulating employees' usage of various generative AI platforms, such as ChatGPT and Google Bard. The primary objective was to prevent inadvertent exposure of sensitive or classified information by scrutinizing the prompts entered by users. Regrettably, this project was ultimately discontinued because it would have been made mostly irrelevant by another project we were assigned to take over that is still in progress.

This project was the development of a full in-house generative AI system. This system is trained using company-specific data and tailored to meet the internal needs of our employees. The overarching goal for this AI system is to serve as an extensive information repository, adept at addressing company-related inquiries, data analytics, policy details, procedural guidelines, and more. Fortunately, the scaffolding of the system was already created; we are utilizing OpenAI's GPT 3.5 model as a base while training it with company data and building a website around it so users could login with their company credentials.

In essence, our technology's primary focus is to resolve the pressing issue of providing employees with a secure, user-friendly, and effective tool for accessing information and generating content in accordance with our company's policies while preventing unintended data disclosures. This initiative underscores the company's new commitment to harnessing AI to enhance internal processes. The new generative AI we are developing not only addresses the immediate challenge of secure information access but also represents a pivot towards company self-sufficiency. By training our AI on proprietary company data, we ensured that it could respond to queries with a level of expertise and specificity that no off-the-shelf solution could provide. Hopefully, it becomes a valuable knowledge resource for employees seeking insights into the company's intricacies.

The Use of Artificial Intelligence in Hiring and Recruitment

For my STS topic, I am examining the topic of fairness in Artificial Intelligence/Machine Learning hiring and recruitment algorithms. Specifically, I want to investigate what potential problems arise in employing AI/ML hiring algorithms in the recruitment process. This is important because AI/ML hiring algorithms are increasingly used by companies to make hiring decisions in lieu of more personal avenues. 24% of companies have adopted artificial intelligence in talent recruitment processes (Regina 1084). From 2018 to 2022, these algorithms' usage has increased by 270%, with over 98% of companies on the Fortune 500 using them (Tilmes 1). Vendors sell companies automated video interview systems, resume filtering algorithms, and even video game tests as applicant screening measures (Andrews & Bucher, 148). Their marketing suggests they are time effective, cost effective, and accurate in their selection for new hires (150). They are even touted by their vendors as an impartial way to screen applicants *without* any form of discrimination (150).

Contrary to the above statement, these algorithms can in fact be biased in several ways, and previous high-profile cases have proven this. For example, the most infamous example is Amazon's attempt at creating an artificial intelligence hiring algorithm was simply a glorified way to discriminate against women (Andrews & Bucher, 147). These algorithms can be trained on data that is biased against certain groups of people or using features that are correlated with race, gender, or other protected characteristics, which will allow them to reflect and even reinforce existing biases (Regina 1091). Additionally, algorithms' opacity and complexity will make them even more difficult to contest or verify (1084). If there is a consistent pattern showing these algorithms are biased and with many more of these algorithms in development, many qualified candidates will be overlooked for jobs, which would have a negative impact on the individuals

unfairly rejected and subsequently on the whole of society. For individuals, being denied a job because of an unfair algorithm can lead to financial hardship, emotional distress, and damage to their self-esteem. For society, unfair AI/ML hiring algorithms can lead to a less diverse workforce, which eliminates perspectives and ideas that stifle innovation and creativity.

To reiterate, my STS research question only focuses on determining if there exists or not a severe problem affecting artificial intelligence and machine learning hiring and recruitment models. This means that the evidence I have gathered will only serve to reinforce the argument that there is a significant problem and lacks any focus on the existence or proposal for a solution or solutions. This is because not only is there significant documentation of the problematic inclusion of bias in most machine learning models (in topics other than hiring and recruitment) regarding any marginalized populations, but there is also no unified solution to this issue. The fact is different models are simply made for different purposes, in different environments, and serving different individuals. Therefore, any specific solution to any potential bias will need to be tailored to that specific model.

The evidence collected includes academic articles and real-world examples of this problem. It consists of meta-analyses, case studies, examinations of existing algorithmic fairness metrics, and proposals for realistic future solutions, whether it be practice or policy. I am looking for patterns in the evidence's assertions regarding any problems in past or present artificial intelligence hiring models but not necessarily any proposed solutions. Naturally, I hope to use that pattern of assertions within the evidence to answer my original research question.

Conclusion

In conclusion, the core focus of my STS research has revolved around the pressing issue of fairness in AI/ML hiring and recruitment algorithms. These algorithms, often considered

unbiased and efficient tools for candidate selection, have been revealed to harbor inherent biases that perpetuate discrimination. The societal implications of this extend beyond individual job applicants, influencing workforce diversity and stifling innovation, which ultimately reverberates through our society.

As I look ahead to my future research paper, I anticipate offering a comprehensive exploration of the extent and implications of bias within AI-driven hiring systems. I will draw upon academic literature, real-world examples, and algorithmic fairness metrics to provide a detailed analysis of the problem. Through the analysis of my evidence, I aim to offer a compelling account of the challenges at hand, focusing on the need for ethical AI development and user-centered design. The expected results of this paper will show the necessity for holistic solutions that consider the broader societal consequences of AI adoption in the realm of hiring and recruitment.

References

- Andrews, L., & Bucher, H. (2022, October 1). AUTOMATING DISCRIMINATION: AI HIRING PRACTICES AND GENDER INEQUALITY. *Cardozo Law Review*, 44(1), 145 - 202.
- Fuchs, L. (2023, January 1). HIRED BY A MACHINE: CAN A NEW YORK CITY LAW ENFORCE ALGORITHMIC FAIRNESS IN HIRING PRACTICES?. *Fordham Journal of Corporate & Financial Law*, 28(1), 185 - 222.
- Hunkenschroer, A. L., & Luetge, C. (2022, July 15). Ethics of AI-Enabled Recruiting and Selection: A Review and Research Agenda. *Journal of Business Ethics*, 178(4), 977 - 1007.
- JONES, C. C. (2022, December 1). SYSTEMATIZING DISCRIMINATION: AI VENDORS & TITLE VII ENFORCEMENT. *University of Pennsylvania Law Review*, 171(1), 235 - 266.
- PESSACH, D., & SHMUELI, E. (2023, April 1). A Review on Fairness in Machine Learning. *ACM Computing Surveys*, 55(3), 1 - 44.
- Raub, M. (2018, June 1). BOTS, BIAS AND BIG DATA: ARTIFICIAL INTELLIGENCE, ALGORITHMIC BIAS AND DISPARATE IMPACT LIABILITY IN HIRING PRACTICES. *Arkansas Law Review (1968-present)*, 71(2), 529 - 570.
- Regina, G. (2023, June 1). DO YOU EVEN KNOW ME?: A.I. AND ITS DISCRIMINATORY EFFECTS IN THE HIRING PROCESS. *Hofstra Law Review*, 51(4), 1081 - 1116.
- Swift, B. (2022, July 1). ARTIFICIAL CONSTRAINTS ON OPPORTUNITY: ARTIFICIAL INTELLIGENCE AND GENDER DISCRIMINATION IN AUTOMATED HIRING PRACTICES FROM AN INFORMATION FIDUCIARY PERSPECTIVE. *Journal of Science & Technology Law*, 28(2), 215 - 245.
- Tilmes, N. (2022, June 1). Disability, fairness, and algorithmic bias in AI recruitment. *Ethics & Information Technology*, 24(2), 1 - 15.

Zeide, E. (2022, December 1). THE SILICON CEILING: HOW ARTIFICIAL INTELLIGENCE
CONSTRUCTS AN INVISIBLE BARRIER TO OPPORTUNITY. UMKC Law Review,
91(2), 403 - 436.