

A Social Construction of Technology Analysis of the Amazon AI Hiring Tool

STS Research Paper
Presented to the Faculty of the
School of Engineering and Applied Science
University of Virginia

By
Rachel Lee
May 11, 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.

Signed: _____
Approved: _____ Date _____
Benjamin J. Laugelli, Lecturer, Department of Engineering and Society

Introduction

In 2014, Amazon collected a group of its engineers to build an automated hiring tool in efforts to recruit top talent with little to no effort on the HR side of Amazon (Goodman, R. 2018). However, this tool incorporated extreme gender bias against women since Amazon engineers were feeding male dominated resume data from the past 10 years into the algorithm (Dastin, J. 2018).

Currently, the failure of the Amazon AI hiring tool is often described as a fault of the technology and can be solved with extensive research and impact facing (Papadopoulos, L. 2018). This viewpoint however fails to shed light upon why this bias occurs within engineers. If we continue to view biased technology a fault of the technology and don't delve deeper into why this bias even exists, we will continue to see companies create gender-biased technology because of the lack of representation of women in tech historically. Analyzing this problem as a deep-rooted societal influence against women in tech will provide a deeper understanding about the technical flaws of the hiring tool.

I will examine the failure of the Amazon AI hiring tool gender bias as a result of society discouraging STEM careers for women from a young age. I will use the Social Construction of Technology theory, which aims to challenge the process of tech design, and that the technologies that result, are socially constructed (Microsoft. n.d.). Applying this concept, I will analyze the process of how the Amazon hiring tool had failed due to social constructs against women in STEM. I will draw upon reports by several sources to support my analysis as there is not firsthand raw data that describes the Amazon hiring tool.

Background

In 2014, Amazon engineers set out to develop a hiring artificial intelligence tool that would automate its recruitment. Its goal was to be handed 100 resumes and pick out the top 5 to hire. ‘The company created 500 computer models to trawl through past candidates’ résumés and pick up on about 50,000 key terms. The system would crawl the web to recommend candidates.’ (Hamilton, I. A. 2018). The hiring tool would rate the candidate on a scale of 5 stars (5 being the best) depending on how closely the candidate resembles a previously successful candidate. Previously successful candidates were the older resumes that were analyzed and parsed for “good” candidate characteristics. However, this included massive amounts of gender bias against women because the AI was combing through mostly male resumes submitted to Amazon over the past 10 years. In a Reuters report, Amazon stated that this tool was never solely used, it was mostly referred to by recruiters (Dastin, J. 2018, October 10).

Literature Review

Many of the analyses surrounding AI hiring tools focus on solely AI ethical issues. These ethical issues are seen as fixable by just being more careful with developing technologies but fail to point out the larger problem as to why these biased technologies were created in the first place.

In An overview of ethical issues in using AI systems in hiring with a case study of Amazon’s AI based hiring tool, Kodiyan analyzes the Amazon AI Hiring Tool through an ethical lens of why the AI failed. Kodiyan states that “hiring is an inherently discriminate process in which some applicant receive offers while other’s don’t base[d] on certain traits that defines a “good” candidate” (Kodiyan, A. A. 2019). The author claims that if hiring is already a biased

process, should having “good” characteristics even be a functional parameter in the hiring technology? Kodiyan argues that it is an implicit trust violation by putting candidates through these hiring technologies. Kodiyan analyzes the Amazon hiring tool in detail with Liffick’s methodology of analysis and even cites a fault to the engineers for not doing enough analysis on the data and considering moral/ethical requirements as functional requirements for the algorithm. His alternative solutions aimed to mitigate AI bias in the hiring space or incorporate some human interaction in the hiring decision. While Kodiyan did a good job in analyzing the engineers’ responsibility and coming up with potential solutions, he didn’t try to raise the question as to why the Amazon engineers didn’t think about the implicit sexist bias the hiring tool would have by using old data with more male resumes being marked as “good.” This implicit bias came from women being deterred from entering tech.

In *Will AI Remove Hiring Bias?*, Lewis highlights that companies with more diverse backgrounds outperform companies that don’t (Lewis, N. 2019, August 16). When analyzing the Amazon hiring tool specifically, she uses an interview with Sarah Brennan, CEO and principal strategist at Accelir, an HR tech consulting firm, to point out that the problem wasn’t from the technology itself, it was the lack of realization that the tool would be using primarily male data being labelled as “good”, since most of the engineers being hired were male (Lewis, N. 2019, August 16). Lewis offers an alternative viewpoint that we can use AI to reduce human bias by using a company called Entelo as an example, where it anonymizes its job candidates (Lewis, N. 2019, August 16). By having humans follow the process, recruiters and HR can be sure that they hired the most qualified candidate (Lewis, N. 2019, August 16). While this article highlights the faults of not knowing or thinking about the biased data in the first place, it also doesn’t attribute the engineers’ biases as a long going issue in our society.

While these literatures discuss very important topics about the ethics of future technology, that being AI, I aim to target the root problem: the reason why the engineers building these technologies are biased. These analyses don't point out the societal biases against women in tech instilled amongst engineers as a viable explanation as to why biased technologies occur. In this paper, I will build on the works of previous scholars by offering an alternative viewpoint that delves into the societal bias against women going into tech. This new viewpoint will offer more of an explanation as to why the Amazon hiring tool failed.

Conceptual Framework

The source of these engineers' biases can be analyzed using the Social Construction of Technology (SCOT) framework. Developed in the 1980s by Trevor Pinch and Wiebe Bijker, SCOT is a framework that views users as the forefront of technological change (Edge et al., 1988), or in other words, that human actions shape technology. It uses the concepts of relevant social groups, interpretive flexibility, closure, and stabilization.

Relevant social groups are the groups who share a meaning of an artifact (Edge et al., 1988). SCOT distinguishes itself from other frameworks in its idea of interpretive flexibility – the claim that different social groups interpret an artifact, or assign meaning to it, in different ways (Mitcham & Johnson, 2005). Interpretive flexibility doesn't continue forever, that is where closure occurs – where so many changes have taken place to the point where it becomes the dominant form of the technology (Edge et al., 1988). This doesn't mean that it's the best of its technology and will result in no rivals. SCOT emphasizes that this process occurs continuously: these technologies may bring up new problems and interpretive flexibility will reappear (Edge et al., 1988).

Social constructivists argue that technological development is influenced by many factors such as social, cultural, economic, and political factors. Engineers look to develop technology that can solve some sort of issue through those factors and those technologies are perceived to have achieved human purpose and improve the social world (Mitcham & Johnson, 2005). They also argue that technology shapes society and society shapes technology (Mitcham & Johnson, 2005). Social forces are much more powerful than we think and have a huge influence on the development and meaning we associate with technology.

The actions of the engineers who developed the Amazon AI Hiring Tool were a product of our societal emphasis that “men are better at STEM.” This exclusion of women in the tech field blinded them to totally disregard the extreme bias incorporated into the AI. I will analyze how these actions have been a result of society’s viewpoint on women in STEM in the following sections.

Analysis

The engineers’ biases behind the Amazon AI Hiring Tool are a societal issue that can be analyzed with the SCOT framework. Society has always encouraged men to go into STEM careers and women to go into non-STEM career paths which has caused a societal bias that has been deterring women from entering the STEM field, and for the purposes of this paper, tech, which I argue was present in the male engineers who developed the Amazon hiring tool.

From a SCOT viewpoint, technologies are socially constructed, therefore, by applying that concept to the failure of the Amazon hiring tool, the negative connotations that have blocked out many women who try to go into tech have socially constructed the Amazon hiring tool and its failure. I will be using the SCOT methodology of interpretive flexibility for my analysis.

The misalignment of interpretation between the developers and women applicants is a reason for the failure of the Amazon hiring tool. The relevant social groups in the design of the hiring tool are the users and developers. I will assign the users as applicants (both men and women) and recruiters and the developers as the engineering team. Men who were applying to Amazon saw the hiring tool as just a hiring tool to get their resumes hopefully shortlisted. However, I believe that the women who were applying to Amazon thought that the hiring tool was flawed and that they might be subject to implicit biases against them. I argue this not only because of firsthand experience as a woman in tech but as a viewpoint that I see amongst other peers as well. The interpretation of the hiring tool on the developers' side aligned with male applicants' interpretation. The interpretation from women applicants showed through obviously when the tool was noticed to prefer men over women. The resumes from previously successful candidates would more likely be males due to the sheer dominance of males in the tech industry which would then lead to an exclusion of a relevant social group (women) during the screening process for this tool. I argue that the engineers' inability to view women as a separate relevant social group than men when developing the hiring tool was a huge factor into the failure of the hiring tool. So why didn't they view women as a separate relevant social group? I believe it is due to the societal construction against women going into tech.

In a UNICEF report mapping gender equality from school to work, there was gender bias present in both parental and teacher expectations in relations to STEM subject fields (Alam, 2020). Parents expected their sons to go into STEM more than their daughters while teachers gave less instructional time and praise to girls in school (Alam, 2020). This data presents a societal expectation or at least more acceptance for boys to go into STEM careers while girls

have been implicitly discriminated in STEM careers. This implicit bias against girls going into STEM related careers translated to the workplace within the male engineers developing the Amazon hiring tool because the gender bias existed for many years before they started to develop the tool.

The gender bias present in the data provided for the ML algorithm shows that the influence of society has a huge impact for the years to come. SCOT cites that society influences the development of technology. Society's viewpoints that discriminate against women in the tech field caused the development of technology to be exclusive because society implicitly supported boys to pursue a career in STEM. This implicit bias from when boys go into school to be treated better or supported more in traditionally STEM classes causes girls to fall behind and be discouraged to go into those fields. By having this socially constructed viewpoint growing up, boys have implicitly picked up the notion against women who go into STEM and tech specifically. When the Amazon engineers were constructed into a team to build a revolutionary AI hiring tool, they saw no problem with using old data from the past 10 years, because they only sought out to solve what they saw as problems through society, the problem that current hiring practices are slow and inefficient. The engineers were blind to underlying society problems because of what problems the engineers only saw or interpreted. This is because of the lack of inclusion in the engineering team to provide a perspective from relevant social groups of the different ways to solve our world's problems.

Some might argue that the data could've been "changed" by using common bias mitigation techniques for machine learning algorithms such as pre-processing and post-processing strategies to reduce the gender bias present. While that is a way to improve the technology and potentially hire more female candidates by reducing the number of downgraded

resumes because of the words those female candidates would commonly use, it wouldn't stop the problem at its root. In the UNICEF report, men and women in the STEM industry were asked about their perception of the reasons why gender diversity was lacking in STEM, men were less likely to attribute the lack of gender diversity to societal discrimination and gender norms (Alam, 2020). This presents a huge problem that I believe was present in the male engineers at Amazon. The inability to think that females might be discriminated to go into tech from an early age resulted in a new piece of technology that reflected that exclusion. I am claiming that the backtracking and understanding of the root problem, the discouragement to women going into tech, has caused the biased hiring tool. Fixing the tool's algorithm would've mitigated the bias temporarily but still wouldn't erase the fact that the data that the algorithm used was extremely in favor of men due to societal norms.

The lack of representation in engineering teams It doesn't matter how technically skilled an engineer is, it's the lack of representation in these engineering teams that cause bias against minorities in tech, most often women. If we hire and collect a diverse set of engineers of all backgrounds, that will naturally create more inclusive technologies.

Conclusion

Many women in tech being steered away or excluded is because of the social construct against women pursuing STEM careers. Using the Social Construction of Technology Framework, the actions of those engineers can be viewed as a result of the societal construct of what an engineer is (males) and thus created a biased. Engineers should always consider all the possible relevant social groups when developing new technology.

References

- Alam, A. (2020, November 20). *Mapping gender equality in STEM from school to work*. UNICEF Office of Global Insight & Policy. Retrieved March 9, 2022, from <https://www.unicef.org/globalinsight/stories/mapping-gender-equality-stem-school-work#:~:text=Women's%20underrepresentation%20in%20STEM%20careers,do%20not%20pursue%20STEM%20jobs>.
- Dastin, J. (2018, October 10). *Amazon scraps secret AI recruiting tool that showed bias against women*. Reuters. Retrieved October 24, 2021, from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.
- Edge, D., Kline, D., & Pinch, T. (1988). Chapter 7: The Social Construction of Technology . In *The social shaping of Technology*. essay, Research Centre for Social Sciences, University of Edinburgh.
- Goodman, R. (2018, October 15). *Why Amazon's Automated Hiring Tool discriminated against women*. American Civil Liberties Union. Retrieved October 24, 2021, from <https://www.aclu.org/blog/womens-rights/womens-rights-workplace/why-amazons-automated-hiring-tool-discriminated-against>.
- Hamilton, I. A. (2018, October 10). *Amazon built an AI tool to hire people but had to shut it down because it was discriminating against women*. Business Insider. Retrieved October 24, 2021, from <https://www.businessinsider.com/amazon-built-ai-to-hire-people-discriminated-against-women-2018->

