

**Gender Bias in Artificial Intelligence: An Analysis of Man-Made Technology and Its
Effects on Women in the Workplace**

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

Mollie Bauer

Spring 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

Advisor

Bryn E. Seabrook, Department of Engineering and Society

Introduction

In *Technology Matters: Questions to Live With*, Nye claims no piece of technology can be fully understood when standing on its own, it must be seen in the context in which it was created. This statement proves true for the emergence of new AI and Big Data practices being used in the job hiring and workplace settings. The important context for this specific piece of technology is the historical discrepancy between men and women in the workforce – in both participation and pay. (Yellen, 2020) Social constructions against women can seep into these technologies and negatively affect progress seen in creating an equal and representative workforce. As such, this paper will explore the following research question: In what ways has the emergence of Big Data and Machine Learning technologies perpetuated marginalization against women in the STEM community?

Research Question and Methods

The paper will address the following research question: In what ways has the emergence of Big Data and Machine Learning technologies perpetuated marginalization against women in the STEM community? To approach this question, the research was guided by a few key phrases that include artificial intelligence, women in STEM, data privacy, Big Tech, and STEM hiring practices. These phrases led to many different sources that discussed complex controversies and case studies. As such, two case studies will be used as methods of analysis. They will be described in detail and be analyzed using the STS frameworks of SCOT and Technological Determinism. After, wicked problem framing will be used to try to dissect the issue of data privacy, and how that ties in with both case studies discussed. The paper aims to paint a clear picture of how this technology has affected women in STEM in recent years as well as help outline a path forward for the STEM industry.

Background

Women in the Workforce

The history of women's participation in the work force is important background to have when analyzing the above research question. In the early 20th century, most women in the United States were not "gainful workers" (White, 2021). Only 20 percent of all women and 5 percent of married women participated in the labor force outside of the home. (Yellen, 2020) These foundational statistics and perpetuated stereotypes are the basis for longstanding disparities between men and women in the workforce. A 2017 poll in the Pew Research Center report found that 50 percent of women said they had experienced gender discrimination at work. The same poll found that only 19 percent of men said the same (White, 2021).

Specific industries today are affected more by these historical disparities. 72 percent of women in the technology industry report being outnumbered by men in meetings by a ratio of 2:1 or more; 26 percent report being outnumbered by 5:1 or more (White, 2021). There is a widespread effort to promote interest of women in engineering and technical roles, but progress to close this gap remains modest (Arefi, 2022). Men have significantly outnumbered women in the workforce throughout history, and specifically in technical roles. This cycle is dangerous and leaves women in a vulnerable position of being left out of technological advancements that will only further prevent them from entering these roles in the future.

Artificial Intelligence and Big Data

Artificial intelligence (AI) is a large field that is used to describe systems that mimic intelligent human behavior to complete complex tasks. Machine learning (ML) is a subset of this

broader field. ML specifically gives computers and systems the ability to learn without explicit programming. Computers can thus program themselves using large amounts of data. Humans choose a ML model that fits the data, and the program finds patterns and makes predictions (Brown, 2021). This means not only can the programmers bias be ingrained in the system, but historical misrepresentations and inequalities can be perpetuated by the system by inputting bad data.

The concept of Big Data is new and only originated in the 1960s. This time period is when large data centers and relational databases were coming to fruition. Big Data refers to large and highly complex data sets that traditional software cannot manage (*What Is Big Data?*, n.d.). The collection of this large amount of data through the emergence of the internet can cause great problems. Personal and confidential information can be gleaned from even the most inconspicuous online engagement or postings. Peoples gender, sexual orientation, and political stances can now be determined through the collection and processing of the massive amounts of online data (Davis, 2017). While this collection of data is not always dangerous or life threatening, the design and use of artificial intelligence models based on these findings and historical prejudice can significantly disadvantage people's lives (Niethammer, 2020).

STS Frameworks

The first main STS framework that will be used to analyze the research question is the social construction of technology (SCOT). SCOT frames the process of technology design as an open process that is collaborated on by different societal groups. The outcomes of that design process can differ based on what groups are present, and there is an assumption that each

subgroup of society will have a say in the way technology is developed. In reality, this assumption is very far from the truth. Historical inequalities amongst races, genders, and socioeconomic levels prevent most subgroups of society from contributing to technology design. The less people who are represented in technology design means the more polarizing and misrepresentative the technology will become (Klein et al., 2002).

This school of thought leads into the second framework that will be used: technological determinism. Technological determinism stresses the power of technology. It has an incredible influence and grasp on society, and effects their processes more than any other known factor. The overarching thought is that technology effects society rather than society effecting technology (Smith, n.d.).

There are two versions of technological determinism: a “soft view” and a “hard view”. The former takes the stance that technological change can drive societal change but can also respond to social pressures. The latter states technological development is completely independent of societal constraints and no pressure can alter its course (Smith, n.d.). In this paper, more often the framework will align with a “soft view”.

Results and Discussion

AI presents three areas of ethical concern: privacy, bias, and the role of human judgment (Writer, 2020). The two case studies examined in this discussion showed deep insight to real world effects of biased technology and the unintended power they hold over certain communities. There are tangible negative effects that these technologies have inflicted against the progression towards equal representation in the STEM fields. These technologies actively

prevented women from applying or getting offered technical positions. The regulation of such technologies is where the three areas of ethical concern come into play and the questions, they pose about responsibility are what makes this problem wicked. Since there is no clear way to regulate AI and cease malpractice in data collection, technology is formed with historical data that carries discriminatory practices. These social constructions have a lasting effect on how AI operates, and due to its precedence and perceived economic importance, is not being overshadowed by present day lobbying and reform.

Facebook Targeted Advertising

According to the Forbes article titled *Facebook Under Fire For Alleged Gender Discrimination In Job Advertisements* a nonprofit group Global Witness released a report that describes an investigation it ran on Facebook's (now Meta) targeted advertisements. The report concluded that Facebook advertisements appeared to operate in a discriminatory manner. The group started this investigation by creating job advertisements using different forms of discriminatory targeting. The four advertisements they used linked to real job postings for mechanics, preschool nurses, pilots, and psychologists. Global Witness chose the objective to deliver ads to the people who are most likely to click on them. As such, 96% and 75% of those shown the ads for mechanic and pilot respectively, were men. On the other hand, 95% and 77% of those shown the ads for preschool nurse and psychologist job postings were women. The report also noted Facebook prompted the ad poster to indicate that the poster will comply with non-discriminatory policies when submitting the ads despite the fact Facebook was not following these policies themselves (Kerpin, n.d.).

This hypocrisy is especially worrying, and as Global Witness states, “Facebook’s business model of profiting from profiling appears to replicate the biases we see in the world, potentially narrowing opportunities for users and preventing progress and equity in the workplace.” Facebook's purposeful perpetuation of these stereotypes only deepens the gap of inequality in not only the workplace, but also society. In recent decades, women have made lots of progress in varying aspects of the job market - from rising employment rates to higher participation in leadership roles in business and government (Geiger, n.d.). The technology is obviously not keeping up with the progress being seen within the American economy as physical demographic changes are being seen but not reflected in things as simple as targeted ads.

Another big issue with a case like this is how unavoidable these targeted ads are on social media sites. Before the digital age, advertisements that were “targeting” certain people were placed in physical places so that that group would see them more. TV commercials would be run depending on the demographic of who watched that program, or maybe jobs for mechanics would have been advertised in men’s magazines. This is far less personal and far more avoidable than curating online ads to each person's digital footprint - all it would take is changing the channel or buying a different magazine that aligned with a specific person's preferences (Joshi, n.d.).

Using the STS frameworks of technological determinism and SCOT, a deeper analysis can be performed on certain aspects of this case. As mentioned previously, despite societal changes in representation, the machine learning technology being used is not perpetuating the trends we are seeing today. This lends to the harder view of technological determinism. Present societal movements to increase representation in the workforce seem to be lost on those at Facebook (*Technological Determinism*, 2015).

The historical social constructions against women in technical fields are having a lingering effect on the technology, though, which is where the SCOT framework becomes important. The social context of this technology helps us understand why the aforementioned inequalities are still around. The historical employment and education gap between men and women and the implicit biases carried by those in industry are a reason the technology is formed in this way. The progress we see in society cannot be reflected in Facebook ads because it is overshadowed by past biased practices grandfathered in (*SCOT*, n.d.).

Using the frameworks to analyze this case, a few points can be concluded. There is an obvious struggle between present societal influences over technology and historical biased social construction. In the end it will take a concerted, and inherently human, effort to change the way technology is presented in today's digital age. The societal progress we make is only obscured when permutations of past biases surface and become unavoidable to a specific audience, as they did in this Facebook Ads case.

Amazon AI Recruitment Tool

The second case is about a faulty AI algorithm used by another tech giant, Amazon. Since 2014 Amazon teams had been building programs to review job applicants resumes in order to streamline and automate the process. As the digital age evolved, automation was a key aspect to businesses' success in the e-commerce industry. The AI tool was used to give applicants scores based on a screen of their resume. The end goal was, “an engine where I’m going to give you 100 resumes, it will spit out the top five, and we’ll hire those” (Dastin, n.d.). The issue that

was noticed in 2015 was that the algorithm was not screening applicants in a fair way - the AI was discriminating against women.

What was uncovered about the program was that the models were trained to screen applicants by observing historical data of resumes submitted and accepted to Amazon in the previous 10 years - most of which were men. As the Reuters article states, Amazon's system, "penalized resumes that included the word 'women's,' as in 'women's chess club captain'" (Dastin, n.d.). People familiar with the system and matter at Amazon also stated the program downgraded or lowered the score of graduates from a few all-women colleges. If these types of clubs and universities had not been well represented in hires at Amazon the past 10 years, it makes logical sense that the program would not associate them with its definition of a successful resume and candidate.

Looking more into the demographic statistics of these types of companies, in the year 2015, female representation at Amazon across all roles reached 39%, but most new minority hires were in the warehouses - not technical and office positions (Rodriguez, n.d.). This is hardly just an Amazon specific problem though. In 2021 across all large technology firms, female representation reached only 33%, up only a few percentage points from the previous year (Hupfer, n.d.). As such, basing resume screening algorithms on historical patterns that have always been inherently unequal was a huge problem, and begs the question why this was implemented in this manner in the first place?

AI has become much more widely used in our daily life and is accepted outside STEM research setting as its original intentions had it. As such, the ethical implications of such a technology are exposed to the world. Michael Sandel, political philosopher states, "Part of the appeal of algorithmic decision-making is that it seems to offer an objective way of overcoming

human subjectivity, bias, and prejudice” (Writer, 2020). The disconnect appears when our assumption of objectivity overshadows obvious design flaws and the subsequent replication of human bias.

Looking at this case through the lens of STS frameworks, one can see again how the historical implications of female discrimination in technical fields has left a lasting impression on our technology today. Society has accepted this technology due to its convenience, apparent objectivity, and growing popularity in many different fields despite the fact it is steeped in past social constructions against a variety of minorities. This also speaks to society's willingness to accept technology that streamlines work and supports capitalistic tendencies, even if it is at the expense of many.

Like the previous case, technological determinism is an apparent framework that can be associated with the effects AI has on people today. This technological tool had the power to prevent subsections of people from getting hired at a job they very well could be qualified for. This is one niche case, but this principle can be applied to a variety of different scenarios in which women have been a minority in.

As Langdon Winner hypothesized, “The technology of a given society is a fundamental influencer of the various ways in which a society exists” (*Technological Determinism*, 2015). This AI technology built upon social constructions and historical biases has the power to affect so many people’s lives for the worse. It was, unintentionally or not, a tool used to perpetuate present day demographics in the STEM field and cloud it by the guise of objectivity.

The Wicked Problem of AI and Data Privacy

An accepted definition of wicked problems is an exceedingly complex problem that lacks clarity in both its aims and solutions, and that is subject to a variety of restraints that prevents multiple attempts at solving it (*What's a Wicked Problem?*, n.d.). The two aspects of this paper and research question, AI regulation and data privacy, can be framed in this way. The Harvard Gazette stated, "Given its power and expected ubiquity, some argue that the use of AI should be tightly regulated" (Writer, 2020). But, as the article goes on to state, there is no consensus or direction on how this regulation should occur and who should make the rules. As of now, most companies that use AI systems are self-policing and rely on reactions from the market and stakeholders to steer their usage and ethics of the systems at hand.

This is where the complexity lies: 1) How do we regulate such technologies to foster more objective systems and a progressive future? 2) Who is accountable if said systems fail and display damaging and discriminatory tendencies? There of course is no clear-cut answer to these problems but they are a great start for a discussion on engineering ethics (Bellamy, n.d.).

It is very widely known, and now widely scrutinized, that companies track people's online movements and collect sensitive information, reflecting in highly personalized advertisements online. Legislation has been created to address these data collection practices by companies, notably the European Union's 2018 General Data Protection Regulation. More controversy came about at the government's limitation of these practices, with people calling legislation like this going too far. There are fears that it will prevent economic growth and western countries like the United States will never be able to keep up with the likes of China, who has a very loose take on data regulation and protection (Candelon, 2021). So, should people

care more about stimulating the economy with the help of data collection or should a person's online privacy have a greater importance?

A different layer to the problem is uncovered when we recognize the link between data privacy and AI. The data being collected in many of these cases is being used by AI algorithms to screen resumes or approve a person for a loan to name just a few. These practices have a very tangible effect on the lives of many, especially women as described through the previous cases. Because of AI's ability to amplify human biases based on how it is trained, it is often wondered if fairness can be coded into the algorithms.

Amazon is again taking a stab at the ethical world of AI as it has developed a new test to ensure fairness in data driven decisions. It is called Conditional Demographic Disparity (CDD) and is defined as being, "the weighted average of demographic disparities for each of the subgroups, with each subgroup disparity weighted in proportion to the number of observations it contains" (Zorio, 2021). The system considers ethical, legal, and technical definitions of "fairness" in an effort to acknowledge, address, and even measure bias within a program. What seems to be forgotten in all of this is that everyone has a different metric for bias and fairness, as seen in previous systems that were made with the intention of being purely objective, and many do not see the legal system as the pinnacle of fairness either.

The problem around data collection and AI regulation is wicked. A big part of the problem is defining what is truly ethical and who can determine this and be responsible for such a system. As it does learn and grow on its own, AI is an intimidating piece of technology, but maybe mostly because it reflects our own rifts and society and finding a solution to the previous solution requires a solution to the latter.

Research Limitations

A limitation of this research is the recency of the cases. They all have surfaced in the last few years, but technology changes quickly and thus could be received as outdated. They were chosen because of the clear way they addressed issues around AI and data privacy as well as how controversial of stories they still are. The research also did not focus much on positive efforts in this field to help protect people's privacy which can paint some companies or industries in a bad light. Further case studies can be performed to help even out the viewpoints and see the bigger picture of things.

Future work on this topic is endless, as the field is so new and ever evolving. I think it is important work for both STS scholars and engineers to pick up as it has implications for everyone. More technical people can benefit from performing research like this because it will only broaden their knowledge and diversify their thought when developing widely used technologies. I think research like this can also serve as a sort of promotional tactic to increase diversity of all kinds in the STEM field.

Conclusion

This research is significant due to its implications on a variety of fields and groups of people. It is only one example of how technology in this new digital age carries unintended pieces of our past. It also serves as a reminder of how powerful unintended bias can be when left unchecked. Women are experiencing unprecedented exclusion from STEM fields despite all the reform being done to diversify the industry. The past social constructions against minorities are being maintained due to the precedent they have set in our economy and world. AI is very much

a reflection of humans but left unchanged it can never learn from its mistakes, only perpetuate what it was taught to be true. It will take a very concerted effort and deep societal reform to change the outlook on data privacy and AI regulations. There is no clear-cut solution, but there is a world in which AI can be optimized to further our society instead of keeping us in the past.

Bibliography

- Arefi, S. M. S. A. (2022). Rethinking Women in Engineering Marginalisation: Beyond Interest Ignition. *2022 IEEE Global Engineering Education Conference (EDUCON)*, 2037–2043. <https://doi.org/10.1109/EDUCON52537.2022.9766589>
- Bellamy, F. (n.d.). *U.S. data privacy laws to enter new era in 2023* / Reuters. Retrieved March 20, 2023, from <https://www.reuters.com/legal/legalindustry/us-data-privacy-laws-enter-new-era-2023-2023-01-12/>
- Brown, S. (2021). *Machine learning, explained* / MIT Sloan. Retrieved September 28, 2022, from <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>
- Candelon, F., Carlo, R. C. di, Bondt, M. D., & Evgeniou, T. (2021, September 1). AI Regulation Is Coming. *Harvard Business Review*. <https://hbr.org/2021/09/ai-regulation-is-coming>
- Dastin, J. (n.d.). *Amazon scraps secret AI recruiting tool that showed bias against women* / Reuters. Retrieved September 27, 2022, from <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>
- Davis, E. (2017, February 15). *The problems of big data, and what to do about them*. World Economic Forum. <https://www.weforum.org/agenda/2017/02/big-data-how-we-can-manage-the-risks/>
- Geiger, A. w, & Parker, K. (n.d.). For Women’s History Month, a look at gender gains – and gaps – in the U.S. *Pew Research Center*. Retrieved March 19, 2023, from <https://www.pewresearch.org/fact-tank/2018/03/15/for-womens-history-month-a-look-at-gender-gains-and-gaps-in-the-u-s/>

- Hupfer, S. (n.d.). *Women in the tech industry: Gaining ground, but facing new headwinds*. Deloitte Insights. Retrieved March 20, 2023, from <https://www2.deloitte.com/us/en/insights/industry/technology/technology-media-and-telecom-predictions/2022/statistics-show-women-in-technology-are-facing-new-headwinds.html>
- Joshi, S. (n.d.). *History of Advertising 101: What You Need to Know*. G2. Retrieved March 19, 2023, from <https://www.g2.com/articles/history-of-advertising>
- Kerpen, C. (n.d.). *Facebook Under Fire For Alleged Gender Discrimination In Job Advertisements*. Forbes. Retrieved October 30, 2022, from <https://www.forbes.com/sites/carriekerpen/2021/09/09/facebook-under-fire-for-alleged-gender-discrimination-in-job-advertisements/>
- Klein, H. K., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, & Human Values*, 27(1), 28–52. <https://doi.org/10.1177/016224390202700102>
- Niethammer, C. (2020). *AI Bias Could Put Women's Lives At Risk—A Challenge For Regulators*. Retrieved September 27, 2022, from <https://www.forbes.com/sites/carmenniethammer/2020/03/02/ai-bias-could-put-womens-lives-at-risk-a-challenge-for-regulators/?sh=86c1455534f2>
- Rodriguez, S. (n.d.). *Amazon 2015 Diversity Report Shows Gains For Minorities, Mostly In Warehouse Roles*. Retrieved March 20, 2023, from <https://www.ibtimes.com/amazon-2015-diversity-report-shows-gains-minorities-mostly-warehouse-roles-2304081>

SCOT / STS Infrastructures. (n.d.). Retrieved March 19, 2023, from

<https://stsinfrastructures.org/content/scot>

Smith, M. R. (n.d.). Technological Determinism in American Culture. . . *The MIT Press*.

Technological Determinism. (2015, March 30). Communication Theory.

<https://www.communicationtheory.org/technological-determinism/>

What Is Big Data? (n.d.). Retrieved February 5, 2023, from <https://www.oracle.com/big->

[data/what-is-big-data/](https://www.oracle.com/big-data/what-is-big-data/)

What's a Wicked Problem? | Wicked Problem. (n.d.). Retrieved March 20, 2023, from

<https://www.stonybrook.edu/commcms/wicked-problem/about/What-is-a-wicked-problem.php>

White, S. (2021, March 8). *Women in tech statistics: The hard truths of an uphill battle*. CIO.

<https://www.cio.com/article/201905/women-in-tech-statistics-the-hard-truths-of-an-uphill-battle.html>

Writer, C. P. H. S. (2020, October 26). Ethical concerns mount as AI takes bigger decision-

making role. *Harvard Gazette*. <https://news.harvard.edu/gazette/story/2020/10/ethical-concerns-mount-as-ai-takes-bigger-decision-making-role/>

Yellen, J. (2020, May 7). The history of women's work and wages and how it has created

success for us all. *Brookings*. <https://www.brookings.edu/essay/the-history-of-womens-work-and-wages-and-how-it-has-created-success-for-us-all/>

Zorio, S. (2021, April 1). *How a paper by three Oxford academics influenced AWS bias and*

explainability software. Amazon Science. <https://www.amazon.science/latest-news/how-a-paper-by-three-oxford-academics-influenced-aws-bias-and-explainability-software>