The Hidden Prejudice of AI: Examining Gender Bias in Algorithmic Decision-Making

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

Grace Weaver

Spring 2025

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

Sean Murray, Associate Professor of STS, Department of Engineering and Society

Introduction

Artificial Intelligence, or AI, has become deeply integrated into many systems of modern society. From personalized recommendations to complex decision-making processes, AI is shaping nearly every aspect of daily life. AI is growing so fast that its influence is expanding across industries at an unprecedented rate, and the Machine Learning market size in the U.S. is expected to grow from $6.49 billion in 2023 to $59.30 billion by 2030 (Growth Analysis, 2023). Revenues from the AI software markets globally have also been steadily increasing, as shown in Figure 1 (Prihatno, Aji & Nurcahyanto, Himawan & Yeong, Min & Jang, Yeong Min, 2020). This rapid expansion highlights the urgency of understanding both the benefits and the potential risks associated with AI systems.

From the job market to law enforcement and beyond, Artificial Intelligence has been trusted to help humans make decisions that are intended to be more efficient and accurate. Consumers often believe that Artificial Intelligence makes decisions that are less susceptible to human error due to human bias (Swift, 2024). This assumption, however, does not account for the ways in which AI can inherit and even magnify existing biases. This belief is incorrect, as real-world applications of AI have demonstrated, often perpetuating and even amplifying societal biases, particularly gender bias.

Figure 1

Global Revenues of the AI Marke in the Last Eight Years

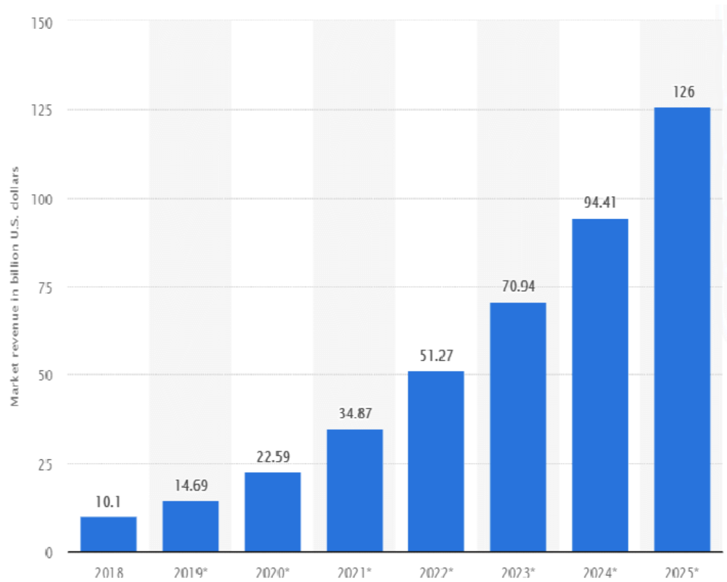

A main concern of AI usage is how, as it is trained on historical and potentially biased data, it can reinforce harmful stereotypes rather than mitigate them. When these biases go unaddressed, AI systems risk becoming tools that sustain and worsen discrimination rather than eliminate it. These biases in AI-driven decision-making can have negative effects on unassuming parties as well. For example, when algorithmic hiring tools were used by Amazon to sort through applications, resumes that included terms more commonly associated with women were rejected. Resumes from male applicants were preferred by the algorithm, and many female applicants were unfairly written off (BBC, 2018). Similarly, facial recognition technologies have demonstrated significant bias, struggling to identify female faces compared to their male counterparts (Singer, N., & Metz, C., 2019). This lack of accuracy can lead to systemic inequalities in employment, security, and other AI-reliant domains. This difference in performance and accuracy can be attributed to the datasets the AI models were trained on and the algorithmic structures used in their development. These biased models can also lead to AI word associations that further exacerbate traditional, and sometimes harmful, gender roles. Studies have revealed that even in seemingly neutral applications, such as natural language processing, AI exhibits gender biases. AI word associations frequently link professions such as “scientist” with men and “nurse” with women (Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A., 2016).

The consequences of untested AI bias are profound. If left unchecked, these biases will continue to shape societal structures in ways that disadvantage certain groups. When hiring algorithms discriminate against women, fewer women secure jobs in technology and leadership roles, perpetuating gender disparities in the workforce (Swift, 2024). In law enforcement, biased facial recognition technology can lead to wrongful convictions. These wrongful arrests can disproportionately target marginalized groups, such as women of color (Singer, N., & Metz, C., 2019). In other basic AI interactions, biased AI-powered recommendations and virtual assistants reinforce harmful stereotypes, which can maintain damaging norms (Harvard Business Review, 2019). The cumulative effect of these biases extends beyond individual cases, contributing to a broader culture of inequality and discrimination. These biases, if left unaddressed, can exacerbate existing inequalities in society rather than positively aid in decision-making.

In this paper, I argue that AI systems reflect and reinforce gender bias due to biased training data, non-diverse development teams, insufficient regulatory oversight, lacking policy enforced by governments, and flaws in the education of the public about AI. These problems can lead to discriminatory outcomes in hiring, law enforcement, and other AI-supported decisions. Addressing these biases requires a multifaceted approach, incorporating insights from social science, technology, and policy studies, and I aim to find the best solutions to the problem.

To support this claim, I will first examine how AI bias manifests in hiring algorithms, illustrating how data-driven recruitment can disadvantage women (Buerk, 2023). I will show how biased facial recognition technologies can perpetuate gender-based disparities, especially in law enforcement. The reinforcement of gender bias through word embedding models will also be analyzed. By assessing multiple facets of the problem, I aim to highlight strategies that can be used to mitigate this bias. Solutions such as ensuring diverse training samples, having diverse development teams, and increasing algorithmic transparency have been proposed. Additionally, I will explore the role of accountability measures and ethical AI guidelines in shaping fairer technological systems. Through my research, I wish to determine the best approach to developing AI systems that promote fairness rather than reinforce discrimination.

Problem Definition

Research has consistently demonstrated that AI systems exhibit gender biases, reflecting and sometimes amplifying societal inequalities. AI bias is observed across various domains, including hiring algorithms, facial recognition technologies, and natural language processing models.

Hiring Algorithms

AI-driven hiring tools have gained popularity in recent years; the CEO of ZipRecruiter estimates that three-quarters of all resumes submitted in job applications are put through a hiring algorithm (Schellemenn, 2022). These algorithms, however, do not exist without problems. A survey found that about 88% of executives knew that their hiring algorithm tools rejected qualified candidates. These algorithms also frequently demonstrate gender bias. Automated systems, such as Amazon's hiring algorithm, learned to favor male candidates over female ones due to biased training data (BBC, 2018). The algorithm was even found to penalize any resume that included the word "women", such as "Women's College" or "Society of Women Engineers" (BBC, 2018). In some cases, algorithms have even been found to automatically reject resumes with career gaps. During the pandemic, more than 54 million women had to leave their jobs to take on caregiving responsibilities, whether that be for elderly parents or young children. Due to this unavoidable gap in employment, women were disproportionately impacted by this AI bias. Gender bias is not the only problem with hiring algorithms. Research indicates that hiring algorithms tend to favor white-associated names over Black male-associated names, with white-associated names being favored 85% of the time (Milne, 2024).

When looking at the current of gender distribution in the corporate world, 46% of entry level roles and only 25% of C-suite and executive roles are held by women globally. There is a clear gender disparity in leadership roles, which already poses a barrier for women trying to obtain jobs. Research also found that men were 33% more likely to receive internal leadership promotions than women (Linkedin). This imbalance is even more prevalent in the tech field, with men holding roughly 80% of executive positions whereas women only hold 20%. These conditions contribute to the bias in AI hiring algorithms. The AI models work with historical data of successful candidates to identify qualified applicants. Many of these datasets contain a majority of male employees, training the models to favor male applicants because they predict they will be more successful (Hall, Ellis, 2023).

Facial Recognition

Facial recognition through AI software has always been a controversial practice. No one consents to have their face used in facial recognition software and it can be used in invasive ways like tracking women going to abortion centers, tracking undocumented immigrants, or targeting people spotted at protests (Fergus, 2024). According to a study from a federal agency, the majority of commercial facial-recognition systems exhibit bias. These algorithms have the highest error rates for African American women. This error is extremely concerning when one acknowledges that false identification can lead to watch-list placements, wrongful arrests, etc. The error rates are the highest for specifically African American women, which is very dangerous considering the criminal justice system already disproportionally targets people of color (Fergus, 2024). There have been calls to have federal agencies stop using facial recognition software. Jay Stanley, a policy analyst at the American Civil Liberties Union, called the technology "dystopian" and demanded the F.B.I. to stop its usage (Singer, N., & Metz, C., 2019).

Research suggests that including participants from more marginalized groups into the data sets that facial recognition AI models are built on would decrease the overall error rates. A study done showed that including more women in a data set of faces significantly decreased the error rate for females. The experiment even included a dataset that had a majority of women's faces, and the error rate for women was lower while the error rate for men did not increase. This again shows that many biases in AI can be caused by a biased or underrepresented data set (Atay, M., Gipson, H., Gwyn, T., & Roy, K.).

Word Embedding

One type of Natural Language Processing (NLP) is word embedding, which represents text as vectors. These vectors can then be used to create relationships between words and predict which words will be used in sequence with other words. Word embedding technology often uses texts from media and popular culture to train its models, and consequently, learns societal biases (Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A., 2016).

When Google's language algorithm is asked to translate gender neutral phrases in different languages into English, there is clear gender bias involved. The gender-neutral phrase "O bir profesör. O bir öğretmen" in Turkish translated to "He's a professor. She is a teacher" in English. Popular language models such as GPT-3 have been shown to associate men with occupations that require higher education and have a higher salary (Caliskan, A., 2021). Similar Natural Language Processing technology has responded in a similar way when prompted to relate occupations to genders. One Model output that "man is to computer programmer as woman is to x" with x=homemaker and that "a father is to a doctor as a mother is to a nurse" (Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A., 2016).

These models are built on many different datasets and sources of language. Often, when these models produce biased text, this text is used for training in other datasets that build other biased language models. There is a hard cycle of biased natural language processing software that is hard to break. This biased language can become present in downstream applications that play a part in tasks such as web searches and job candidate selection (Caliskan, A., 2021).

The Gap in Action

Despite increased awareness, solutions for mitigating AI gender bias remain insufficiently implemented. Many recommendations have been proposed, but adoption across industries is inconsistent. Scholars emphasize the need for diverse training datasets, yet many AI systems continue to rely on historically skewed data (del Villar, Z., 2025). Calls for algorithmic audits and transparency in AI decision-making processes have grown, but regulatory measures are still lacking (McNew, D. 2020). Government action, such as prohibiting purchases of biased AI systems, has been suggested as a necessary step toward reducing bias in law enforcement applications (Fergus, 2024). The growing reliance on AI in hiring, law enforcement, and everyday decision-making makes addressing gender bias an urgent priority. Future research should explore more effective bias mitigation strategies and advocate for industry-wide standards to ensure equitable AI systems.

Methods

Mitigating gender bias in AI requires a multifaceted sociotechnical research approach that examines the issue from historical, technical, and societal perspectives. By integrating multiple analytical frameworks, this study explores the complexity of AI bias and the most effective strategies for addressing it. Through historical analysis, system analysis, and expert insights, this research aims to uncover both the root causes of bias and actionable solutions for a more equitable AI landscape.

Historical analysis will be used to research the history of AI bias and see how it has shaped decisions and technology. STS professor Kathryn Neeley states that historical analysis in sociotechnical research “organizes evidence in chronological order and goes beyond simple narrative construction to identify both continuity and change over time" (Neeley, 2024,). It would be interesting to see how AI has changed and developed and what policies have been enacted to help moderate any negative effects. This will give the reader a better understanding of the issue and its historical significance. The knowledge of AI's history will also help frame for the reader how widespread an affect AI bias can have and introduce the reader to any actors that are at play in these technologies.

I will also use system analysis throughout my research. System analysis is used to identify how systems interact and affect societal structures, human values, and technology. Neeley describes systems analysis as focusing "on relationships between interrelated components that form a whole" (Neeley, 2024). System analysis will help me consider the large actor network in AI. Actor Network Theory, or ANT, aims to recognize and show the connections between different actors in a societal network. Bruno Latour, a key developer of Actor Network Theory, describes actors as "Something that acts or to which activity is granted by others. It implies no special motivation of human individual actors, nor of humans in general. An actant can literally be anything provided it is granted to be the source of an action" (Latour, B., 2017). There are many people and objects that interact and have a stake in the AI decision making process; these are known as "actors" in this network. It is important to acknowledge and study the effects each actor can have on one another. The development teams and the data used to train the AI models are both actors that have a lot of impact on the output of AI. These decisions then impact the groups of women who are the victims of the gender bias ingrained in this AI. Figure 2 below shows how different actors in this network can impact each other.

Figure 2

Actor Network of AI and its Bias

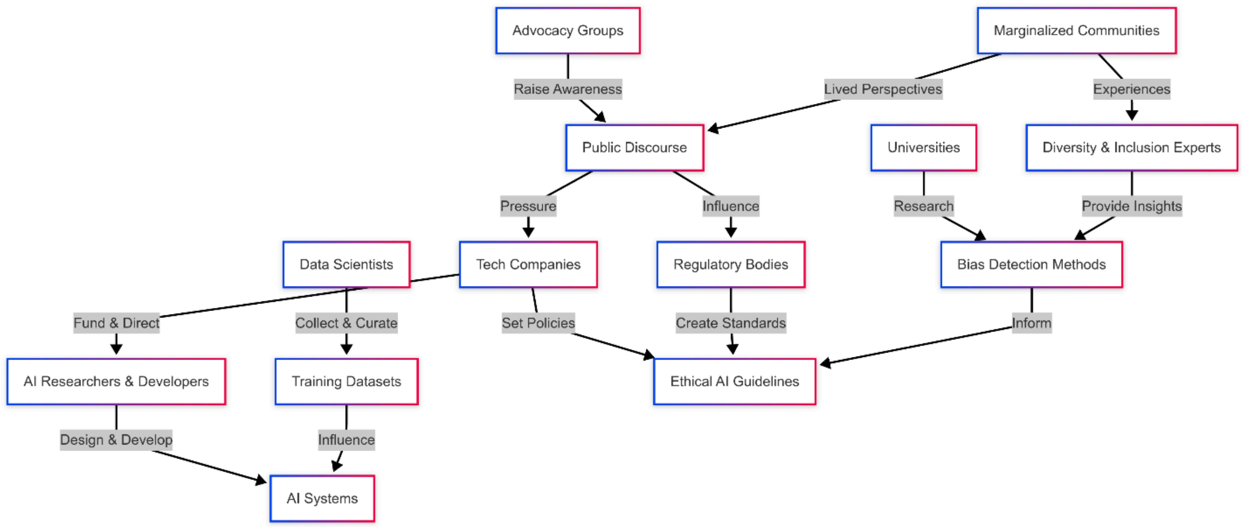

By employing both historical and system analysis, this research will provide a comprehensive view of how AI gender bias emerges, persists, and affects different societal actors. Understanding the historical trajectory of AI bias will help contextualize the issue, while system analysis will reveal the intricate relationships between technology, developers, and affected individuals. These combined approaches will ensure a well-rounded exploration of AI gender bias, ultimately contributing to more effective and informed strategies for its mitigation.

In addition to historical and system analysis, this research incorporates insights from AI ethics experts, researchers, and policymakers who have extensively studied algorithmic bias and its mitigation. Understanding expert perspectives provides a critical foundation for identifying both the root causes of AI gender bias and the most effective solutions proposed to date.

To gather expert opinions, I analyzed academic papers, reports, and interviews from AI researchers specializing in algorithmic fairness, ethics, and policy. These experts provided key insights into how bias emerges in machine learning models, the limitations of current bias mitigation techniques, and the ethical challenges in addressing AI discrimination.

Beyond expert perspectives, this research also examines empirical studies that demonstrate instances of AI bias in real-world applications. Case studies of biased hiring algorithms, gender disparities in AI-generated content, and discriminatory facial recognition systems serve as concrete evidence of the issue. These studies not only highlight the prevalence of bias but also showcase technical and policy-driven solutions, such as bias-aware model training, fairness constraints in machine learning, and standardized auditing practices.

By integrating expert opinions with case study findings, this research develops a well-rounded approach to bias mitigation. The combination of theoretical frameworks, empirical evidence, and expert insights ensures that the proposed solutions are both informed by research and applicable in real-world AI development.

Results

Addressing gender bias in AI requires a multifaceted approach that includes improving data collection, diversifying AI development teams, implementing strong regulatory policy, and increasing transparency around the inner workings of AI.

AI systems learn from training data, and when this training data contains societal bias and stereotypes, the resulting decision model will inevitably reinforce those biases (del Villar, Z., 2025). A study on AI test grading shows that having a balanced gender distribution in training data leads to fairer decisions (Latif, Zhai, Liu). Another study on multiple common facial recognition systems found that increasing the amount of female faces in the training data significantly improved the accuracy rate in identifying females. The male accuracy rate did not decrease when more female data points were added to the training data, showing that there is no drawback in including more females (Atay, M., Gipson, H., Gwyn, T., & Roy, K.). Developers must be actively selecting data that includes minority groups and reflects many cultural backgrounds to train their models (del Villar, Z., 2025).

In 2024, women made up only one third of employees in the United States tech industry (United States; Bureau of Labor Statistics; Lightcast; 2024). Due to this, most of the AI decision-making technologies development teams are made up of men. Many stereotypes and biases that women would be aware of and have experienced, would go unnoticed during testing. One would think that promoting more women to work in the tech industry could eliminate instances of AI gender bias, but there are other barriers preventing women's participation. Globally, about 50% of women in STEM fields experience sexual harassment in their jobs(Manasi, Panchanadeswaran, Sours, 2023). There must be more support structures in place to promote and aid women that are already in STEM fields. If the field becomes a more welcoming place for women, there would be a more equal distribution of gender. This would lead to more perspectives on development teams, and more biases that are recognized and removed.

There has been an increase in support for AI bias regulation and policy in the last few years. Experts suggest different ideas on what the most effective policies would be to effectively decrease instances of AI gender bias. Some claim that the US congress should pass a disparate impact law that covers the impact of AI bias. This would allow users of AI to sue companies for bias they are victims of without needing to prove that a decision maker intended to discriminate against them. This would be applicable to AI bias since users are not aware of the inner workings behind AI decisions. This possibility of legal action against them would promote companies to more thoroughly test their AI models before deploying them to the public (Bains, C., 2024). Others say that the governments should enact laws and guidelines for AI development that companies must follow. The European Union has already proposed the EU AI Act which requires AI systems to go through intense bias testing and compliance tests. There are also several international documents such as the ISO/IEC TR 24027:2021 and IEEE P7003 that describe how AI should be developed to avoid bias. These documents should help guide developers while they are building their AI models. In New York City, any Automated Employment Decision Tools (AEDTs) must be submitted to an independent third-part bias audit every year. Moving forward, the development of policy and regulation of AI testing for bias will ensure that AI can benefit everyone and promote a more equitable society (Byrne, Lee, Le, 2024).

Often AI does not deliver decisions with uncertainty; AI seems confident and authoritative. It is because of this that humans seem to blindly trust AI decisions. They often believe that decisions from AI are more accurate than human decisions. It is important that users are educated on the presence of bias in AI. They must understand that AI can produce incorrect or skewed information, and they should always question AI decisions. Informing the public about the problems within AI development and bias could lead to less malicious activity against those that AI is programmed to biased against (Busiek, J., 2024).

AI gender bias is a multifaceted issue that requires proactive intervention at every stage of AI development, deployment, and regulation. By prioritizing diverse training datasets, fostering inclusive AI development teams, enforcing legal accountability, adopting ethical standards, and educating the public, AI systems can become more equitable and just. As AI continues to shape decision-making in critical areas such as hiring, loan approvals, and law enforcement, these mitigation strategies will be essential in preventing the reinforcement of societal inequalities.

Conclusion

AI, while often perceived as objective and unbiased, is deeply influenced by human decisions at every stage of its development and deployment. As a result, it clearly reflects and amplifies societal biases—particularly gender bias—due to flawed training data, non-diverse development teams, and inadequate oversight. The examples of biased hiring algorithms, facial recognition errors, and discriminatory word embeddings illustrate the profound impact of these biases, shaping critical decisions in employment, law enforcement, and everyday AI interactions. These biases are not just technical flaws but ethical concerns that have far-reaching consequences for individuals and communities worldwide.

Addressing AI gender bias requires an intersectional and multidimensional approach. Ensuring diverse and representative training datasets is a crucial first step in mitigating bias, as studies have shown that increasing representation can lead to fairer AI outcomes without sacrificing accuracy. Additionally, increasing diversity within AI development teams is essential, as more inclusive perspectives can help identify and challenge biases that might otherwise go unnoticed. Without these varied perspectives, AI models risk being trained in environments that reflect only a narrow segment of society, leading to exclusionary and even harmful outcomes. Beyond these technical solutions, stronger regulatory frameworks, such as bias audits and compliance laws, must be enacted to hold companies accountable for the consequences of their AI systems. Governments, industry leaders, and academic institutions must work together to establish clear guidelines and enforcement mechanisms to ensure AI is developed and deployed responsibly.

Equally important is public awareness and education about AI bias. The perception of AI as infallible must be challenged so that users of AI-driven technologies critically evaluate AI-driven decisions rather than accepting them at face value. Educational initiatives, media coverage, and advocacy efforts can play a significant role in shifting public understanding and encouraging demand for ethical AI practices. As AI becomes increasingly integrated into society, failing to address its biases will only reinforce and perpetuate existing inequalities. However, with intentional, ongoing efforts in research, policy, and ethical AI development, we can ensure that AI serves as a tool for fairness rather than discrimination. By fostering a culture of accountability, transparency, and inclusivity in AI, we can shape a future where technology empowers rather than marginalizes, ensuring that AI contributes to a more just and equitable society.

Works Cited

4 ways to address gender bias in ai. (2019, November 20). Harvard Business Review. https://hbr.org/2019/11/4-ways-to-address-gender-bias-in-ai

Adamovic, Mladen, and Andreas Leibbrandt. "A Large-scale Field Experiment on Occupational Gender Segregation and Hiring Discrimination." Industrial Relations: A Journal of Economy and Society, vol. 62, no. 1, Jan. 2023, pp. 34–59. DOI.org (Crossref), https://doi.org/10.1111/irel.12318.

Atay, M., Gipson, H., Gwyn, T., & Roy, K. (n.d.). Evaluation of Gender Bias in Facial Recognition with Traditional Machine Learning Algorithms. https://par.nsf.gov/servlets/purl/10358566

Bains, C. (2024, September 13). The legal doctrine that will be key to preventing AI discrimination. Brookings. https://www.brookings.edu/articles/the-legal-doctrine-that-will-be-key-to-preventing-ai-discrimination/

BBC. (2018, October 10). Amazon scrapped "sexist AI" tool. BBC News. https://www.bbc.com/news/technology-45809919

Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A. (2016, July 21). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. arXiv.org. https://arxiv.org/abs/1607.06520

Buerk, Emily. Technological Connection: Improving Resource Access through Technology; How Does AI Contribute to Existing Bias in the Technological Hiring Process? 2023. University of Virginia. DOI.org (Datacite), https://doi.org/10.18130/N59Z-KA49.

Busiek, J. (2024, March 21). Three fixes for AI's bias problem. University of California. https://www.universityofcalifornia.edu/news/three-fixes-ais-bias-problem

Byrne, A., Lee, D., & Le, Q. (2024, October 14). Bias in AI: Tackling the issues through regulations and standards. Public Policy. https://publicpolicy.ie/papers/bias-in-ai-tackling-the-issues-through-regulations-and-standards/

Caliskan, A. (2021, May 10). Detecting and mitigating bias in Natural Language Processing. Brookings. https://www.brookings.edu/articles/detecting-and-mitigating-bias-in-natural-language-processing/

del Villar, Z. (2025, February 5). How AI reinforces gender bias-and what we can do about it. UN Women – Headquarters. https://www.unwomen.org/en/news-stories/interview/2025/02/how-ai-reinforces-gender-bias-and-what-we-can-do-about-it#:~:text=%E2%80%9CVoice%20assistants%20defaulting%20to%20female,%22scientist%22%20with%20men.%E2%80%9D

Fergus, R. (2024, February 29). Biased technology: The Automated Discrimination of Facial Recognition. ACLU of Minnesota. https://www.aclu-mn.org/en/news/biased-technology-automated-discrimination-facial-recognition#:~:text=Studies%20show%20that%20facial%20recognition,Facial%20recognition%20automates%20discrimination.

Kubiak, Emeric, et al. “Gender Equity in Hiring: Examining the Effectiveness of a Personality-Based Algorithm.” Frontiers in Psychology, vol. 14, Aug. 2023. Frontiers, https://doi.org/10.3389/fpsyg.2023.1219865.

Latif, E., Zhai, X., & Liu, L. (n.d.). AI Gender Bias, Disparities, and Fairness: Does Training Data Matter? https://arxiv.org/html/2312.10833v2

Latour, B. (2017). On actor-network theory. A few clarifications, plus more than a few complications. Philosophical Literary Journal Logos, 27(1), 173–197. https://doi.org/10.22394/0869-5377-2017-1-173-197

LinkedIn. "Gender Equity in the Workplace." LinkedIn, https://linkedin.github.io/gender-equity-2022. Accessed 30 Mar. 2025.

Manasi, A., Panchanadeswaran, S., & Sours, E. (2023, March 17). Addressing gender bias to achieve ethical AI. IPI Global Observatory. https://theglobalobservatory.org/2023/03/gender-bias-ethical-artificial-intelligence/

McNew, D. (2020, May 1). The problem of bias in facial recognition: Strategic Technologies Blog. CSIS. https://www.csis.org/blogs/strategic-technologies-blog/problem-bias-facial-recognition

Milne, S. (2024, October). AI tools show biases in ranking job applicants' names according to perceived race and gender. UW News. https://www.washington.edu/news/2024/10/31/ai-bias-resume-screening-race-gender/

Neeley, Kathryn. (2024). Sociotechnical Research: What? Why? How? Who?

Prihatno, Aji & Nurcahyanto, Himawan & Yeong, Min & Jang, Yeong Min. (2020). Artificial Intelligence Platform Based for Smart Factory. 10.1109/DASC/PiCom/DataCom/CyberSciTe.

Schellmann, H. (2022, May 11). Finding it hard to get a new job? robot recruiters might be to blame. The Guardian. https://amp.theguardian.com/us-news/2022/may/11/artitifical-intelligence-job-applications-screen-robot-recruiters

Singer, N., & Metz, C. (2019, December 19). Many facial-recognition systems are biased, says U.S. study. The New York Times. https://www.nytimes.com/2019/12/19/technology/facial-recognition-bias.html

Swift, J. (2024, July 22). Algorithmic bias in job hiring. Gender Policy Report. https://genderpolicyreport.umn.edu/algorithmic-bias-in-job-hiring/

Tran, Justin. Developing Wearable Headband for Enhancing Slow Wave Sleep in Older Adults; Bias and Discrimination in Hiring Algorithms. 2024. University of Virginia. DOI.org (Datacite), https://doi.org/10.18130/JZKA-TR75.

Whitney, M. (2022, June 7). Is AI bias in hiring software keeping women from the workforce?. LinkedIn. https://www.linkedin.com/pulse/biased-ai-keeping-women-from-workforce-michele-whitney/