How Technology Shapes Gender Bias

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

Daniel Knorr
Spring, 2021

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

Signature ____ Date __5/5/2021________
Daniel Knorr

Approved ____ Date __5/4/2021____
Hannah Rogers, Department of Engineering and Society

Abstract

Gender bias in the development and application of digital technologies leads to unequal outcomes for users of such products. The issue of gender bias is not unique to digital technologies; it is something that has been shaping technology development of all kinds. Prior to 1993, the majority of medical studies were limited to male participants. Researchers treated men to be the "norm" of the population and assumed women would respond the same as men to treatments. Since then, increased diversity in the groups involved with medical research has been correlated with better patient outcomes. A similar reckoning is at work in the broader technology industry where male dominated firms continue to perpetuate gender biases, regardless of intent. Through unrepresentative teams, a desire for ever faster growth, and the notion that machines present no bias, gender inequities in technology persist. By examining cases of biased technologies as political artifacts and approaching the problem from a techno-feminist perspective, major deficiencies in hiring practice and design theory are found to be largely responsible. Wajcman's feminist theories of technology provide guidelines for addressing the politics of sexist artifacts. By following a few guidelines, technology developers may introduce products that better serve all of society. By having diverse teams of decision-makers, fairer hiring practices, and independent audits of artificial intelligence-based applications, there can be better outcomes in digital technology for all users.

How Technology Shapes Gender Bias

"Revolutions in technology do not create new societies, but they do change the terms in which social, political and economic relationships are played out" (Wajcman, 2006). The growth of digital technologies provides an opportunity for gender biases to play out on a massive digital scale. Without a change in how digital technologies are designed and applied, they will continue to be shaped by and perpetuate gender bias. A similar crisis took place in the medical technology industry. Headaches, nausea, bleeding, and seizures are common adverse drug reactions that occur for women who are prescribed some FDA approved medications. The same reactions are not common in men who are prescribed the same medications. Similarly, digital technology (referred to as "technology" in this paper) and the development thereof does not treat men and women equally. Until 1993, there was no requirement in the FDA approval process for clinical trials to report the outcome by sex or include both sexes as a covariate (Zucker & Prendergast, 2020). Before 1993, the overwhelmingly male researchers, doctors, and scientists who developed medications did not include women in clinical trials. For one, "males, frequently of the Caucasian race, were considered to be the "norm" study population" (Liu & DiPietro Mager, 2016). By the 1980s, a U.S. Public Health Service task force found "that the quality of knowledge related to women's health was lacking due to the exclusion of women in research" (Liu & DiPietro Mager, 2016). In the years since the passage of the US National Institute of Health (NIH) Revitalization Act of 1993, the medical community has been forced to address the varied effects of prescription medications on the basis of sex. It is time for a reckoning of similar proportions on the varied effects of digital technologies on the basis of sex as well. Let us examine how technology shapes gender bias and how gender bias shapes technology.

Langdon Winner's framework of political technologies distinguishes between two sets of political technologies, those which are "technical arrangements as forms of order" and those which are "inherently political" (1980). When considering the relationship between recent technologies and the propagation of gender biases, it follows to examine it from the case of a "form of order" (1980). This is because "technological change expresses a panoply of human motives, not the least of which is the desire of some to have dominion over others" (1980). The goal of this examination of technologically backed gender bias is to uncover ways in which gender inequality is rooted in technology. Since inequality suggests one group having more power than another, the idea of dominion expressed by Winner lends his framework to the dissection of the issue.

Practices Which Drive Biased Development

Just as medications had variable effects on patients of different genders, there are other technologies that do not serve users of different genders equally. Where there were physical inequities in medication, there are algorithmic inequities in technology products. Consider the case of the Xbox Kinect where "the system had been tested on men ages 18 to 35 and didn't recognize as well the body motions of women and children" (Tugend, 2019). The Xbox Kinect is another example of men designing for themselves and leaving women out of the process, and thus rendering them underserved by the technology. Not all cases of gender-biased technologies are as harmless as the Xbox Kinect. Amazon employed a resume screening algorithm that discriminated against female applicants. Law enforcement agencies in Florida employed a technology developed by a third party that regularly allotted longer jail stays to female inmates than they normally should have received (Hamilton, 2019). The overall trend seems to be that supposedly unbiased machines and algorithms are actually perpetuating the biases of their creators, be they explicit or not. The preeminent feminist STS writer Judy Wajcman noted that "gender is connected to other axes of power such as race, colonialism, sexuality, disability, and class" (2010). Addressing the relationship of technology to all of these areas would be a monumental task. However, it is the notion of intersectionality that makes understanding the relationship between biased technologies and gender important. By employing Langdon Winner's framework of political technology, it will become clearer how these technologies "embody specific forms of power and authority" (1980). The use of Judy Wajcman's feminist theories of technology will help propose remedies for the gender biases rampant in the technology industry. Together, these frameworks allow a closer examination of how technology challenges or reinforces gender biases. With any luck, the technology industry may soon experience a reckoning on gender discrimination similar to that felt by the medical community only a couple decades ago.

The dominion of men over women in technology is not hard to uncover. Without even considering the differences in power of the individual, consider the differences in power of the many. In the early days of computing, programming was considered as a woman's work domain (Wachter-Boettcher, 2018). Early graduating classes from computer science departments were more balanced in gender than they are today (National Center for Education Statistics, 2012). However, after the introduction of the personal computer and of its marketing primarily towards men, programming took on a gendered role. By making the technology development space a place dedicated to men, a type of observer bias was introduced. Like medical researchers who did not expect women to react differently to medications, developers would find themselves creating biased solutions whether they knew it or not. Thus, by sheer numbers alone, men can claim dominance in the development of technology. The imbalance gets worse as one moves up the ladder of seniority in technology firms. In senior management positions at Facebook, for example, 66% of posts are occupied by men (Williams, 2020). Thus overall, and especially where it matters, women find themselves the overwhelmed minority in the development of technology. Winner suggests that power begets power and that technologies are used to further the form of order that their creators desire. Thus, we see the worsening of the gender disparity in computing that started in the 1980s and continues today.

In the case of work place technologies, this perpetuation of power is evident in the case of the Amazon hiring algorithm that was shut down in 2018. In this example, a resume screening technology used machine learning to sift out the resumes that were most similar to successful previous hires (Lavanchy, 2018). The model was trained using past hiring data from recruiting executed by humans. At first glance, this seems to be an improvement. Removing people from the process would seem to be the best way to make the hiring decisions fairer. If women were just as qualified as men, then more would surely advance through the pipeline. This was not the case. Since the model was trained using data from past recruiters, the data it was provided with was inherently biased. The model was trained on data that embodied the biases of the human hiring managers. Consider instead if Amazon had noticed these biases built into the system. They could have taken steps to test statistically for adverse impact. This is known as the adverse impact ratio. "Originating in the employment context in the 1970s, the ratio consists of dividing the proportion of the selected group in the disadvantaged class by the proportion of selected members of the advantaged group" (Burt, 2020). This ratio was developed by the U.S. Department of Labor with the Equal Employment Opportunity Commission and adopted as part of the Uniform Guidelines on Employee Selection Procedures. Generally, if the ratio is under

80%, there is strong evidence for discrimination. Tests like these are possible and only discovering disparate impact in deployment should not be permissible.

The issue with the Amazon hiring algorithm is not unique. It is also very likely the engineers did not have any malicious intent. The problem is that without diverse teams, designers miss out on pain points that do not relate to their own experiences with technology. Consider the story of Fatima who worked on a team designing a new smart watch. Her superiors wrongly assumed that women would prioritize style over function in such a device (Wachter-Boettcher, 2019). Fatima brought up her concerns to the relevant decision makers, but was ignored. Here, the predominantly male team did not even seriously consider the concerns Fatima brought up. She was able to draw on her own experiences and the issues she noticed in the product were exactly what caused it to be a massive failure. If Fatima's team could not create a gender equitable product with one woman present, there is little hope for teams composed exclusively of men. The path to better technology that serves users fairly, then, is through diverse, gender-representative teams. While it could be possible to achieve better outcomes without changing team composition, such as testing for adverse impact, the issue of observer bias remains. As was discussed previously in regards to clinical research, observer bias can have long lasting effects and is best removed from the development space.

A major difference between the case of the Amazon hiring algorithm and Fatima's story is that in the latter, there was a woman present on the design team who could uncover the source of the bias directly. In the other case, the implicit biases of the algorithm were only uncovered after it had been put to use. The issue with only discovering bias after technology is deployed is that sometimes it cannot be uncovered. The obsession with human centric design (HCD) is partly to blame. HCD evangelist Don Norman spearheaded the idea that the user should be abstracted

away as far as possible from the technology. Things should just work for people. While this is a nice sentiment, there is an issue with his claim that “it is not our duty to understand the arbitrary, meaningless dictates of machines” (1988). This notion that “dictates” of machines are “arbitrary” and “meaningless” ignores the reality that machines embody the socializations of their designers. As Winner suggests, these dictates have political will and they are not meaningless. When machines are making important decisions on behalf of humans with real life implications, it might be critical to understand the underlying technology. In the case of the hiring algorithm, knowing that it was trained on biased data
completely changes whether it is seen as an actual improvement in hiring practice or not.

Issues in hiring like the problem presented previously create a vicious cycle. There is a misrepresentation of gender in technology development which is followed by sexist technologies (likely not even intentionally) which are used to hire more of the same people which leads to a greater gender imbalance. Without concerted effort, the problem will not go away.

Thus far it has been established how the development of technology leads to ingrained biases against women. Furthermore, popular theories of design then worsen the issue by hiding the sexist technicalities from users. If these are the vehicles of power, then, in accordance with Winner’s framework, to what end are they used? I propose they are used to further entrench those in power. Let us first examine the language of the technology industry- specifically the existence of job titles like “Innovation Evangelist,” “Software Ninjaneer,” and “Full Stack Magician” (CB Insights, 2019). The use of such titles hides the true meaning of the work behind them. They are a way of convincing outsiders that the creators of the technology possess some innate special quality that they do not. While some of these job titles are more obviously fantastical than others, the very term ‘software engineer’ is problematic. “Traditional engineers are regulated, certified, and subject to apprenticeship and continuing education" (Bogost, 2015). In the modern era of continuous integration and deployment (CI/CD), software engineers are able to circumvent many of the usual quality controls that a civil engineer would have to follow. A bridge must work the first time, but the technology industry has deemed it acceptable for software to roll out iteratively. The now-retired Facebook motto of "move fast and break things" highlights this difference between software engineering and other engineering disciplines (Taneja, 2019). One would be hard pressed to find a non-tech engineering firm that would be comfortable living under such a creed. Like the previously discussed job titles, the term 'software engineer' is itself a method of entrenching power. With this language, developers attempt to convince the public that they are of the same class as civil and public works engineers and deserve the same trust.

The Language of Biased Development

Why does this language matter from a gender perspective? Recall the previously discussed trend of reduced participation of women in the software engineering field. 'Programming' was initially considered work suitable for women. Female 'computers' of the Apollo missions later became NASA's first programmers going into the early 1970s (Atkinson, 2014). But, by the 1980s, 'software engineering' was increasingly dominated by men. This coincides with the NATO Software Engineering Conferences of 1968 and 1969. Sponsored by major defense contractors, these summits sought to make software development more similar in discipline to other engineering fields (Randell, 1996). Evidenced by the industry's modern reliance on CI/CD, those efforts seem to have failed. 61% of the 3,650 developers surveyed in the GitLab "Mapping the DevSecOps Landscape" report affirmed deploying a product multiple times per day (2020). What those conferences did accomplish, though, was the cementing of the term ‘software engineer’ in American corporate parlance. By switching from ‘programmers’ to ‘software engineers,’ the once female-friendly field adopted the existing sex biases associated with other engineering disciplines. “Male computer programmers sought to increase the prestige of their field, through creating professional associations, through erecting educational requirements for programming careers, and through discouraging the hiring of women” (Frink, 2011). There was a chance for computing to be free from the gender discrimination rampant in other engineering fields, but those hopes were dashed when the field morphed into our modern conception of software engineering.

Technology itself is primed to continue this language-based discrimination against women. Word embedding is a common method for natural language processing. In this framework, text is represented as vectors that can then be used in machine learning applications. According to a 2016 study, “word embeddings trained on Google News articles exhibit female/male gender stereotypes to a disturbing extent” (Bolukbasi et al., 2016). The system correctly correlated “man” and “woman” with “king” and “queen.” However, when provided with “man”, “woman,” and “computer programmer,” the trained word embeddings complete the relationship with “homemaker.” Like the hiring algorithm before, these word embeddings take on the bias of the data they were trained on. One might not see using news data to train language processing as political. If the word embeddings discovered a relationship between “woman” and “homemaker” that was as strong as the relationship between “man” and “computer programmer,” perhaps the system has just correctly identified a truth about the current status of society. The system is representing the status quo. Inaction, however, is political. Support for the status quo, is still political. Just because women work as homemakers more frequently than they do as computer programmers (when sifting through news), does not mean algorithms should be designed to keep it that way. Cavalier disregard for gender bias in word embeddings is how technology like the previously discussed hiring algorithm ends up changing lives for the worse.

The Financing of Biased Development

So far, we have examined how the composition of teams can lead to biased technology. The technology used in hiring and the language of the industry then contribute to the imbalance of genders on development teams, thus consolidating power. Designers and CI/CD methods were also assigned some blame, but the power of financing cannot be ignored, and venture capitalists too have a role in perpetuating gender bias in technology. How has this group perpetuated the development of biased technology? How have they used the politics of technological artifacts to their advantage? Let us first consider the flow of capital in the industry. According to data from PitchBook, only 2% of capital raised in 2017 went to companies founded by women (2020). Either women were discriminated against in the allocation of funding, or there were not enough startups founded by women available to receive funding. Either way, there is a problem. If founders tend to be overwhelmingly male, then the previously described problems of implicit bias persist. Without mixed composition in a team, as shown in Fatima's story, then there are no fresh perspectives. Just as this held true for design teams, it holds true for management. The contribution of Venture Capitalists to sexism in technology does not end with their role in the allocation of funds. In 2017, 'more than two dozen women in the technology start-up industry spoke to The Times in recent days about being sexually harassed" (Benner, 2017). This is another example of the abuse of power by entrenched actors in the technology space. Prior examples, like the hiring case, highlighted procedural abuses. In this case, the abuse is personal.

By keeping women from founding their own companies, either by lack of funding or harassment, the technology industry succeeds in undermining the interests of women in the technology development process.

The Perpetuation of Bias

The importance of technology development led by women cannot be understated. In her seminal work *Feminism Confronts Technology*, author Judy Wajcman discusses how the male-led development of contraceptive technologies resulted in technologies that policed women's bodies and primarily served the interests of men. For example, contraceptive pills with considerable side effects have been made available for mass market use by women. However, male based approaches with similar or fewer side effects were never approved for widespread use. The onus and health burden of family planning was firmly placed on women. Without women's agency in the development of software, women may too be left with the pernicious side effects of male developed artificially intelligent systems. One of these side effects is the cementing of the notion of women working in support roles for men.

One of the most common manifestations of this phenomenon occurs in voice assistant technologies. "In the United States, Siri, Alexa, Cortana, and Google Assistant—which collectively total an estimated 92.4% of U.S. market share for smartphone assistants—have traditionally featured female-sounding voices" (Chin & Robison, 2020). This choice relates back to the societal expectation that women be caring and helpful more so than this is expected of men (Parker, Horowitz, & Stepler, 2017). The issue with this default is made more pressing when one considers the generation of children growing up as native users of virtual assistants. Female-voiced bots are "a powerful socialization tool that teaches us about the role of women, girls, and people who are gendered female to respond on demand" (Lever, 2018). Consider the political and social signals such devices send. Those who do not identify as female are conditioned to view those who do as service individuals. According to a UNESCO report, "the prominence of female-sounding voice assistants encourages stereotypes of women as submissive and compliant" (Chin & Robison, 2020). In general, the male dominated teams that were discussed previously regularly reinforce their primacy via virtualized women working service jobs. There is one famous example where male consumers rejected a female assistant and preferred a male voice. According to the research of Stanford University Professor Clifford Nass, German automaker BMW was forced to recall an early version of their voice navigation system that featured a female voice. It was found that German male drivers did not like taking directions from a female assistant (Nass & Yen, 2010). While this example may seem to contradict earlier claims, it does not. In this case the male voice was preferred because the bot was issuing commands, not simply serving up requested information. The common thread is that assistive bots have cemented for many the idea that women should be helpful and men be authoritative.

Steps Towards Bias Reduction

If so many problems arise in technology development as a result of male dominated design, management, and venture capital teams, then what does the alternative look like? Are women-led organizations more aware of implicit biases that can be built into technology? The recent success of Pymetrics in this space would suggest that it is possible. The Pymetrics story highlights how firms with women in agentive positions can develop technology that combats gender biases. Pymetrics was co-founded by Dr. Frida Polli in 2013. As of April 2021, Polli still serves as the company's chief executive officer. The company has received almost $60 million in funding (Crunchbase, 2021). According to the company's mission statement, their candidate assessment technology "leverages new science and audited technology to improve prediction and reduce bias" (Pymetrics, 2021). Recruiters for a leading global investment firm used the Pymetrics platform for the prescreening of job applicants and ended up with a 62% increase in female representation. Recruiters were 95% satisfied with the candidates advanced by the AI platform (Pymetrics, 2021). A cooperative audit conducted by researchers from Northeastern University found that pre-employment assessments developed and deployed by Pymetrics met fairness criteria set forth by the Uniform Guidelines on Employee Selection Procedures as dictated by the United States Office of Federal Contract Compliance Programs. Furthermore, no bias appeared as a result of adverse impact testing in regards to the treatment of the seven protected demographic categories outlined by the United States Equal Employment Opportunity Commission (Wilson et al., 2021). Pymetrics accomplished this by carefully designing their data ingestion practices and regularly testing their product for disparate impact.

There are a few important lessons in the Pymetrics case. The first is that the development of bias-reduced (or even free) technology is possible, even in cases as notoriously bias ridden as machine learning. On this first point, Winner would suggest that all technologies are still inherently political. Perhaps bias cannot be totally removed, but at least technology can be made to better reflect the values of the general population, instead of the values of the designers. The second is that women's perspectives are pivotal in the development of these technologies. Pymetrics did not set out initially to address machine learning hiring bias. In an interview with her alma mater, Harvard Business School, Dr. Polli reflected on becoming a tech entrepreneur that "it was a shock to [her] to go to meetings or dinners or whatever the case may be and just constantly be surrounded by only men" (Polli, 2019). It was her experience as a female founder that led her to tackle this mammoth issue. The third important aspect of this case to note is that Pymetrics should not be the norm. Carefully auditing algorithms for disparate impact should be, but placing the onus of building fairer technologies on female founders should not be. Building fairer technologies starts with having people of varying backgrounds in decision-making roles. Furthermore, just as development is continuous, adverse impact testing can be continuously done after release as well. Achieving equitable outcomes in the development and application of technology products should be an active and vigorous activity of all industry participants. To that end, Pymetrics "has provided their adverse impact testing framework as an open source tool" (Wilson et al., 2021). By providing their audited tool to other companies in their technology space, Pymetrics is making good on their mission to reduce algorithmic hiring bias in all products, not just their own.

How can engineers just starting out in their careers follow the example set forth by engineers like those who built Pymetrics? The combating of gender bias in technology must occur on two fronts. Engineers must prevent their own biases from shaping the development of technology, and be sure that what they build does not perpetuate such biases in others. The former can be achieved by working in teams of a balanced gender composition. More importantly, this includes management teams and decision makers, lest a similar situation to Fatima's arise. The third-party auditing of a new technology, such as the kind undertaken by Pymetrics, is an excellent option for validating a final product. The latter can be achieved by not hiding too much from the user. Allow them, before first use, to choose their own assistant voice, for example, to avoid the malicious power of the default option. Where Norman suggests hiding as much as possible from the user, do the opposite and inform the user how their actions and thoughts are being guided.

More women in technology leadership roles in existing firms may correlate to more startups being founded by women. Perhaps this will begin to make a dent in the gross gender imbalance in funding provided by venture capitalists. Additionally, it might be time to change job titles in the industry to better reflect the skills they require. Overall, it is the language, design paradigms, hiring practices, and funding decisions of the industry that allow gender bias to shape the development of technology. It is a lack of due diligence with artificially intelligent systems and a poor choice in defaults that allow technology to shape gender biases. There is hope to change course in the technology industry, as has been exemplified by Pymetrics. The importance of unbiased digital technologies will only grow in the years after the Covid-19 Pandemic. As more and more jobs are made remote, digital technologies will play a larger role in governing human interactions. If the biases present in the use and development of digital technologies today are not addressed, then there will be many, not only women, left behind in the digital economy of the future.

References

Atkinson, J. (2015, August 24). From computers to Leaders: Women at NASA Langley. Retrieved April 06, 2021, from https://www.nasa.gov/larc/from-computers-to-leaders-women-at-nasa-langley

Benner, K. (2017, June 30). Women in Tech Speak Frankly on Culture of Harassment. The New York Times. Retrieved April 6, 2021, from https://www.nytimes.com/2017/06/30/technology/women-entrepreneurs-speak-out-sexual-harassment.html

Bolukbasi, T., Chang, K., Zou, J., Saligrama, V., Kalai, A. (2016, July 21). Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embedding (Tech.). Retrieved April 6, 2021, from Conference on Neural Information Processing Systems website: https://papers.nips.cc/paper/2016/file/a486cd07e4ac3d270571622f4f316ec5-Paper.pdf

Burt, A. (2020, October 08). How to fight discrimination in ai. Retrieved April 06, 2021, from https://hbr.org/2020/08/how-to-fight-discrimination-in-ai

CB Insights. (2020, June 26). The 25 most absurd job titles in tech. Retrieved April 06, 2021, from https://www.cbinsights.com/research/most-absurd-tech-job-titles/?utm_source=angellist

Chin, C., & Robison, M. (2020, November 23). How AI bots and voice assistants reinforce gender bias (Rep.). Retrieved April 6, 2021, from The Brookings Institution website: https://www.brookings.edu/research/how-ai-bots-and-voice-assistants-reinforce-gender-bias/

Crunchbase. (2021, March 10). Pymetrics - Crunchbase company profile & funding. Retrieved April 06, 2021, from https://www.crunchbase.com/organization/pymetrics

Frink, B. D. (2011, June 1). Researcher reveals how "computer Geeks" replaced "Computer Girls". Retrieved April 06, 2021, from https://gender.stanford.edu/news-publications/gender-news/researcher-reveals-how-computer-geeks-replaced-computer-girls

GitLab, Inc. (2020). Mapping the DevSecOps Landscape (Rep.). Retrieved April 6, 2021, from GitLab Inc. website: https://about.gitlab.com/developer-survey/#download

Hamilton, M. (2019). The Sexist Algorithm. *Behavioral Sciences & the Law*, 145-157. doi:https://doi.org/10.1002/bsl.2406

Lavanchy, M. (2018, November 15). Why Amazon's sexist AI recruiting tool is better than a human. Retrieved April 06, 2021, from https://www.imd.org/research-knowledge/articles/amazons-sexist-hiring-algorithm-could-still-be-better-than-a-human/

Lever, E. (2018, April 26). I Was a Human Siri. Intelligencer. Retrieved April 6, 2021, from https://nymag.com/intelligencer/smarthome/i-was-a-human-siri-french-virtual-assistant.html

Liu, K. A., & DiPietro Mager, N. A. (2016). Women's involvement in clinical trials: Historical perspective and future implications. Pharmacy Practice, 14(1), 708-708. doi:10.18549/pharmpract.2016.01.708

Nass, C. I., Yen, C. (2010). The man who lied to his laptop: What machines teach us about human relationships. Prince Frederick, MD: Recorded Books.

Parker, K., Horowitz, J. M., & Stepler, R. (2017, December 5). On Gender Differences, No Consensus on Nature vs. Nurture (Rep.). Retrieved April 20, 2021, from Pew Research Center website: https://www.pewresearch.org/social-trends/2017/12/05/on-gender-differences-no-consensus-on-nature-vs-nurture/

PitchBook. (2020, December 7). All In: Female Founders and CEOs in the US VC Ecosystem (Rep.). Retrieved April 6, 2021, from PitchBook website: https://pitchbook.com/news/reports/2020-all-in-female-founders-and-ceos-in-the-us-vc-ecosystem?utm_medium=nl-na&utm_source=reports&utm_campaign=2020-all-in-female-founders-and-ceos-in-the-us-vc-ecosystem

Polli, F. (2020, October 08). Using AI to eliminate bias from hiring. Retrieved April 06, 2021, from https://hbr.org/2019/10/using-ai-to-eliminate-bias-from-hiring

Pymetrics. (2021, April 5). Talent matching platform. Retrieved April 06, 2021, from https://www.pymetrics.ai/

Randell, B. (1996). Dagstuhl Seminar 9635. In Dagstuhl-Seminar 9635: "History of Software Engineering" Schloss Dagstuhl, August 26 - 30, 1996. University of Newcastle upon Tyne.

Taneja, H. (2019, January 23). The era of "move fast and break things" is over. Retrieved April 06, 2021, from https://hbr.org/2019/01/the-era-of-move-fast-and-break-things-is-over

Tugend, A. (2019, June 17). Exposing the Bias Embedded in Tech. *The New York Times*. Retrieved November 30, 2020, from https://www.nytimes.com/2019/06/17/business/artificial-intelligence-bias-tech.html

United States, National Center for Education Statistics. (2012). Digest of Education Statistics. National Center for Education Statistics.

Wachter-Boettcher, S. (2018). Technically wrong: Sexist apps, biased algorithms, and other threats of toxic tech. New York, NY: W.W. Norton & Company.

Wajcman, J. (2006). *TechnoFeminism*. Cambridge, UK: Polity Press.

Wajcman, J. (2010). Feminist Theories of Technology. *Handbook of Science and Technology Studies,* 189-204. doi:10.4135/9781412990127.n9

Williams, M. (2020). Facebook 2020 Diversity Report: Advancing Opportunity for All (Publication). Menlo Park, California: Facebook.

Wilson, C., Ghosh, A., Jiang, S., Mislove, A., Baker, L., Szary, J., . . . Polli, F. (2021, March 1). Building and Auditing Fair Algorithms: A Case Study in Candidate Screening (Rep.). Retrieved April 6, 2021, from Northeastern University website: https://dl.acm.org/doi/abs/10.1145/3442188.3445928

Winner, L. (1980). Do Artifacts Have Politics? *Daedalus*, *109*(1,), 121–136.

Zucker, I., & Prendergast, B. J. (2020). Sex differences in pharmacokinetics predict adverse drug reactions in women. *Biology of Sex Differences,* (11). doi:10.1186/s13293-020-00308-5