



Impact of Biases in Artificial Intelligence on the Job Market

Akash Erra

Abstract

As artificial intelligence has become more prevalent in an ever-advancing society, so have AI based programs in various sectors of the job market. Historically, advancements in automation and technology have been seen to bring profound change to the structures of the job market and society, and artificial intelligence can be grouped in that same category of technological advancement. With AI being used to a higher degree in hiring processes, tech jobs, and automation, there is a rising fear about increased unemployment rates, economic inequality, and it being far more difficult to get certain jobs. On the other hand, artificial intelligence may allow for more jobs in other fields and can allow for enhanced productivity and growth in various business sectors. Despite its advancement, artificial intelligence still exhibits imperfections—particularly its potential for bias in the job market. For example, artificial intelligence hiring algorithms can exacerbate existing biases drawn from past data and can cause unfair hiring practices against any given race, gender, age, or other metric. Along with unfair hiring, biases in AI can lead to wage inequality/unfair employee evaluation, and on top of all of that, AI biases can have an effect on other parts of society, such as healthcare and criminal detection, by targeting certain races, genders, or age groups using existing data. While AI may posit some benefits (enhanced productivity, possibly more jobs), there still may be more downsides that follow (biases, increased unemployment rates, economic inequality).

Introduction

As artificial intelligence (AI) has become more prevalent in our ever-advancing society, so have AI-based programs become prevalent in the job market for automating tasks, sorting through large pools of data, and overall increasing efficiency while reducing costs. This paper explores what the net impact of AI could be on the job market, and crucially, how bias in AI can interact with the impact. Historically, advancements in automation and technology have brought profound changes to the structures of the job market and society; artificial intelligence can be grouped in that same category of technological advancement. With AI being used to a higher degree in hiring processes, tech jobs, and automation, there is a rising fear about increased unemployment rates, economic inequality, and it being far more difficult to get certain jobs. On the other hand, artificial intelligence may create more jobs in other fields and can allow for enhanced productivity and growth in various sectors.

While artificial intelligence is very advanced, it still has a significant imperfection: its biases. Biases in artificial intelligence can have a serious effect on the job market. For example, artificial intelligence hiring algorithms can exacerbate existing biases drawn from past data and can cause unfair hiring practices against any given race, gender, age, or other demographic. Along with unfair hiring, biases in AI can lead to wage inequality/unfair employee evaluation; AI can also exacerbate economic inequality by displacing workers and creating skill mismatches, leading to more disparities in both job opportunities and wages. Additionally, they can have an effect on other parts of society, such as healthcare and criminal detection, by targeting certain races, genders, or age groups using existing data. While AI may posit some benefits, such as increasing productivity and potentially increasing demand for certain jobs, there still may be more downsides that follow, like increased unemployment rates in certain industries, increased bias influence, and increased economic inequality.

In this paper I describe the theory on how AI may affect unemployment rates and historical patterns from similar technologies, I analyze common biases in current AI systems and the impacts they have had on different communities, and I conclude with a discussion of the potential impact from the combination of these.

How Artificial Intelligence Affects Unemployment Rates

Artificial intelligence has had varied impacts on unemployment and the broader job market, and experts on the topic have different perspectives on what the true effect of artificial intelligence on the job market will be (Frank et. al, 2019). First, it is important to understand what is most likely to come from the introduction of artificial intelligence on the job market: jobs will be both created and destroyed.

Throughout world history, automation and technological advancements have been recurring trends, which often correlate with job losses as automation increases (Bessen, 2020). In the context of the historical trend, while these job losses do occur, it's not entirely caused by automation. In an example based on US textile production (Bessen, 2020), it can be shown that the number of US textile production workers has fallen since the 1940s mainly due to automation, but about 100 years before 1940, the employment rate of workers was growing despite automation. So why did employment rates fall since then, even though they were keeping up earlier? Consumer demand and industry growth may have played a significant role. As more textile goods were being produced due to automation, prices went down, and consumer demand increased. Because of the large increase in demand, the need for employment was high, and therefore, more workers in the textile industry were being hired even as each worker was producing more.

Job losses primarily resulted from declining demand caused by an oversupply of textiles; with reduced demand, businesses scaled back production and, consequently, employment. Artificial intelligence (AI) will likely have similar impacts, as AI falls under the same category of automation.

Today, AI is used for automation in manufacturing, data analysis and organization, marketing, customer service, and even in hiring job applicants, and the conclusion drawn in this example can be applied to artificial intelligence. As with previous automation waves, AI adoption may reduce jobs by taking over repetitive tasks across industries. However, much like the textile industry other industries saw forms of growth that outpaced the job loss after automation; manufacturing industries needed positions for people to manage automated systems, healthcare staff to overview AI's processing of medical data, data scientists interpreting data synthesized by AI models, etc., essentially, wherever AI is, there will likely be a need for someone to overview it, and in-turn industries will need people to fill that role of overviewing AI.

The idea that jobs can both increase and decrease after automation is known as "creative destruction", a term popularized by Joseph Schumpeter. Creative destruction refers to the process of new innovations and technologies making old ones obsolete, leading to the destruction of existing economic structures while creating new job opportunities at the same time. For example, looking back to the U.S. textile industries as well as steel and automobile industries, employment has been steadily declining for years due to automation and globalization (Grennes 2003). Despite these losses in this specific sector of the economy, the whole of the U.S. saw economic growth and job growth; for each textile job lost, 36 more jobs in other industries were created (Grennes 2003). While this same pattern of net job creation may continue today, its effects are unbalanced, as modern automation and AI tend to benefit certain

industries and communities more than others. Creative destruction isn't limited to jobs being destroyed in one industry but being created in another; creative destruction can also happen within industries. For example, the software engineering industry may undertake significant change due to AI automation, causing a loss in jobs for entry-level software engineers and reducing the demand for routine tasks, but there will also likely be an increase in jobs offered for managing AI systems—effectively carrying out creative destruction by destroying existing job positions while simultaneously creating new ones.

With more and more businesses using AI to a greater degree than before, AI implementation may also generate new job opportunities. In a paper by Wilson et al. (2017), the authors suggest new positions like being a “trainer, explainer, and sustainer” to be seen in the job market. A “trainer” is a job that essentially teaches artificial intelligence systems how they should be doing their jobs. For example, an AI customer service chatbot will have to be trained by a human in order to gain the ability to mimic human-like behaviors. An “explainer” is a job that connects the technological aspects of AI with the business world, explaining and connecting the outputs to decisions or decision makers. Their ability to explain the concepts of their AI system and the trustworthiness of their AI systems effectively to business leaders will be extremely important, as an increasing number of businesses are embracing AI. Lastly, a “sustainer” is a job that oversees AI automation and makes sure that it is carrying out its job both properly and effectively. They are responsible for maintaining the AI, verifying the outputs, and modifying the AI over time. All the types of jobs listed previously are already performed and will likely increase as businesses will only advance in artificial-intelligence related automation, and they need real people to work in positions like this as these types of jobs are likely to not be automated in the near future. While these new types of jobs will be created, it is unclear how many there will be, as the number of AI systems “trainers”, “explainers”, and “sustainers” can each work on may also increase due to new technologies.

Although AI creates some jobs, it will likely reduce overall employment. A way to convey this definite loss in jobs will be most explicitly seen in the comparison of low-skill versus high-skill jobs. Low-skill jobs generally don't require much specialized training and are usually repetitive tasks that can be learned quickly. Conversely, high-skill jobs require specialized training and qualifications towards a specific position, and usually involve non-repetitive tasks and creative thinking, including devising new tasks based on the position's responsibilities.

The “repetitive tasks” component of low-skill jobs demonstrates why these jobs suffer in comparison to high-skill jobs when confronting artificial intelligence implementation. AI is primarily introduced to increase efficiency, especially in automating simple, repetitive tasks. For example, cashier work involves repetitive tasks that machines can easily replicate, hence why around 96% of retail stores in the U.S. have self-checkout machines; they increase efficiency and do the repetitive tasks better than a human could, and generally, these machines are much better at doing these types of tasks because they do not experience fatigue like real people do.

Now, in the context of high-skill jobs, these cannot be easily replicated by an artificial intelligence system. Non-repetitive tasks can be difficult for AI systems to perform well in, as they are generally trained for certain outputs based on the inputs they have previously seen. When faced with unfamiliar data or circumstances, the AI may fail to appropriately adapt. Additionally, jobs like doctors and lawyers require a certain level of human qualities and behaviors like judgment, empathy, and ethical/moral values. Jobs that require creative thinking and human-specific qualities are more difficult to automate, and may not be so easily replicated by AI.

However, just because high-skill positions don't have easily adaptable artificial intelligence implementations, it doesn't mean that AI can't take over other jobs in the same sector. For instance, in the lawyer example given previously, legal researchers could potentially lose out on jobs due to superior document analysis performed by artificial intelligence machines (Koenig et al., 2019). Similarly, AI has been used in medical contexts to diagnose health conditions, analyze medical images and data, and predict patient outcomes with comparable accuracy to human doctors (Tai, 2020). These examples demonstrate that AI can still have an impact on the number of jobs in certain professions or businesses, even if it can't directly replace the high-skill jobs. Furthermore, as AI capabilities advance, the distinction between low-skilled and high-skilled jobs from the computer's perspective may become less clear.

Skill mismatches are a likely byproduct that will come from the implementation of AI. As the capabilities of AI develop further and as AI is applied to an increasing number of jobs, people will struggle to adapt to the new skill ceiling required to take on certain tasks, and therefore are more likely to lose their jobs and effectively be displaced by AI. This phenomenon creates structural unemployment by means of creative destruction, which is defined earlier as the replacement of existing economic structures, or jobs, by new innovative processes.

The skill mismatches of today, specifically the ones caused by AI, aren't the same as the ones caused by previous forms of automation, because AI is a lot more advanced than previous forms of automation. AI can do tasks that require high-level decision making, pattern recognition, and problem-solving, whereas past automation mainly involved machines that took over simple, repetitive physical tasks, which were subsequently easier for people to adapt to. This creates a larger gap between the replaced skills and the new skills needed to stay relevant. The main issue is that people will have a hard time adapting to technological skill-shifts required by the implementation of AI, especially as AI becomes more and more capable.

Figure 1

Demand for Skills in the Context of Automation

Skill Attributes	Example of Skills	Projected Fluctuations in Skill Demand Due to Automation	Impact on Jobs
Physical and Manual Skills	Operating vehicles, stocking and packaging products, energy and mining	Decrease in the United States and Europe	Jobs that require mainly physical and manual skills are declining due to automation
Technological / Basic Digital Skills	Advanced IT / programming	Increase in the United States and Europe	Jobs that require technological skills will rapidly grow in the future, along with automation.
Social and Emotional Skills	Management, Leadership positions, Communications	Increase in the United States and Europe	Jobs that require social and emotional skills will rapidly grow despite the increase of automation in the future.
High Cognitive Skills	Creativity, critical thinking, decision making	Increase in the United States and Europe	Jobs that require higher cognitive skills are surging along with an increase in automation.

Note. All fluctuations in demand are projected within the years 2016-2030. Demand fluctuation statistics from Bughin et al., 2018.

Using low-skill versus high-skill jobs to convey the potential job losses in the job market is just one example, but there are still more ways that AI can come for people's jobs. Rather than using terms as "low-skill" and "high-skill" jobs, we can take a look at the attributes of specific skills that are being demanded in the job market and group these into physical/manual, technological/digital, social/emotional, and high cognitive skills. Figure 1 shows the specific fluctuations in the demand for these skills in the context of rapid automation in the future.

Looking at jobs that require physical and manual skills, there will likely be a decrease, as has been the trend for the past 15-20 years (Bughin et al., 2018). Being that automation and advanced technology usually shine in fields that require physical/manual labor due to the likelihood of repetitive tasks, this trend comes as no surprise, as the world has been constantly advancing in this aspect. While certain physical tasks are non-repetitive and may require

creative adaptation, as the motor capabilities and ability for AI systems to learn increase, so does the likelihood of automation for those more difficult physical tasks.

Along with this increase in advanced technology will come the need for more workers that are experienced in tech-related fields, and because of that need, we can see the increased demand for all technological skills, even basic digital skills. Jobs in the technology sector or those driving innovation generally require a high level of education, whereas jobs that mainly deal with physical labor do not have that same requirement.

The other skills in the table, social and emotional skills as well as high cognitive skills, fall under a different category in the job market, as they can't truly be connected to being "high-skill" or a "low-skill" position. In this case, it is more of the fact that these types of jobs cannot be easily replicated—or cannot be done at all—by artificial intelligence or any other automation. Looking at social and emotional skills and in the specific context of artificial intelligence, it is very much obvious to see why these skills cannot be replicated: it is extremely difficult for the objective nature of AI to capture the full range of the subjective nature of human emotion. AI technology specializes in reading and understanding objective data, so naturally, AI systems would be well versed in mimicking human emotion or comprehending emotion through surveys, social media, and consumer data. At a glance, it may seem that by this fact, AI systems can perceive human emotion—which is possible to an extent—but solely relying on data to interpret human emotion comes with its drawbacks: one of them being bias. In a study by Lauren Rhue (2019), facial recognition AI technology found itself trying to interpret human emotion in facial expressions, but the technology showed significant bias when it came to black faces: the technology labeled these faces as angrier than white faces, and associated black men's facial expressions with threatening behavior. Because of reasons like bias, it is simply both difficult and unreliable for AI to be tasked with capturing human emotion, with today's technology at least.

In contrast, cognitive skills present a different case. Automation and specifically artificial intelligence can potentially replace basic cognitive skills such as basic data entry, routine problem-solving, and pattern recognition, which is why the projected fluctuation of jobs that require only basic cognitive skills are declining (Bughin et al. 2018), and why specifically high cognitive skills are projected to have an increase in demand.

Using the examples and descriptions provided above, it is easy to notice that the implementation of artificial intelligence/automation into businesses is not as straightforward as it seems. Potentially, more AI can lead to job losses, causing less of a need for low-skill and easily automated tasks, but on the other hand AI can lead to an increase in the potential number of jobs by fostering new AI-related job positions or jobs that cannot be easily replaced by AI.

Biases in Artificial Intelligence

While artificial intelligence is used in many complex ways for all types of jobs, specifically to improve performance and productivity, it can also bring up new challenges, such as introducing a new factor of bias. Artificial intelligence systems are not flawless and can be biased due to the data they are trained on. Being that some AI systems are trained on historical datasets that contain human cognitive biases or may suffer from non-random sampling, it is possible that, without any intention of the programmer, the AI is biased, which can be detrimental in many ways

A notable example of AI bias occurred at Amazon (Weissman, 2018), where a hiring algorithm favored male applicants over female ones. The AI system was likely biased towards

male applicants due to the past training data used on the AI. The training data was likely conducted by humans, and therefore had its own hidden biases in the data that was inherited by the AI system, leading to the preference of male applicants over female ones. Another example that reflects the same type of bias can be seen in gender biases with career STEM advertisements when both were using “unbiased” algorithms (Lambrecht and Tucker, 2019).

While the most common examples of biases usually lie in gender, racial biases have also been observed. An example of racial bias can be seen in US criminal software, known as COMPAS, which decides whether or not to release an offender. An investigation into the software found that it was more likely to assign higher risk detection scores to African-American’s compared to caucasians with the same otherwise similar profiles (Dressel and Farid, 2018). Again, this bias was likely perpetuated by the past training data used in the AI system, where the humans who took the data had their own biases which were subsequently inflicted onto the AI. This directly displays how biases in AI systems can have a clear and direct impact on not only someone’s job, but a person’s life; people are staying in prison longer because of biases in artificial intelligence.

Essentially, there are many cases that have to do with bias in artificial intelligence, and these can have varied and important impacts on the job market. For example, drawing back to the Amazon situation, there is a clear and direct impact of bias affecting a certain group of people, in this case, it is specifically women. It’s clear that the algorithm’s bias marginalized an entire group, and in doing so, this situation clearly displays a negative effect of artificial intelligence on the job market. The main issue at hand and possibly the most direct way that bias by AI can negatively affect the job market is through the use of AI-based hiring algorithms. The use of such algorithms by firms and companies is for the purpose of efficiency, and while it is a lot easier to sift through job applications and pick through the most qualified applicants using an algorithm, it comes with its drawbacks: marginalizing entire groups of people. A substantial amount of large and mid-size businesses use AI-based hiring algorithms frequently, which can serve to convey the true depth behind the issue. When looking at the context of Amazon’s situation, the past data used to train the AI were mainly resumes by men, and therefore the AI model took a significant bias against applicants that were male. The AI model would put a red flag on applicants with keywords like “women” or any other female-related activity on their resume. The predominance of male resumes in the training data likely reflects existing human biases in hiring practices, and the fact that AI models were trained on this past data goes to show how these hiring algorithms can perpetuate human biases. Because of this, it is extremely difficult to fully eliminate biases from any AI-based algorithm, as most of the past data algorithms are trained on data collected by humans.

Although outside the job market, the earlier example of racial bias in criminal justice further illustrates the persistence of algorithmic bias, when US criminal software was targeting African-Americans with the same profiles as caucasians. The AI-based software was trained on past data used by humans, and therefore the biases that the human data held were effectively given to the software, which is reflected in its targeting of a certain group of people. What is more significant to realize is that a large amount of data is biased, because humans are innately biased.

Whether it is through explicit biases that people are consciously aware of, or through implicit biases that are hard to detect, humans carry bias in some type of way. Now, while data can be cleaned to dispel the impact of bias, it cannot be fully eradicated, only mitigated. The fact that bias can never be dispelled reflects on AI, and it serves the same fact that biases in some

AI systems can only be reduced, not removed. The impact of this on the job market is detrimental, as many businesses may have AI-based hiring algorithms that may carry biases towards specific applicants and may contribute to the marginalization of certain groups of people due to preexisting structural inequalities.

Artificial intelligence can affect jobs beyond just the hiring process, too. Some examples of this can be seen when artificial intelligence is implemented into human resources operations. Artificial intelligence that is used for performance reviews can possibly favor certain groups of people and disenfranchise others, effectively undervaluing certain groups and affecting their job status, denying career advancement. Another example can be seen in wage discrimination, where employees' wages are adjusted to offer lower wages to groups that have been underrepresented. AI-bias can also prohibit people from even seeing certain job opportunities or gaining any insightful connections. For example, platforms such as LinkedIn may have biased AI models that display certain job opportunities over higher-value ones based on user profile information, limiting the opportunities of certain groups of people.

Companies that use AI models to hire people or are in any way related to jobs could be negatively affected. Just for using these models, these companies risk having a damaged reputation and spending a great deal of money to “fix” these AI systems. For example, let's say an AI algorithm was used to hire prospective workers and review their resumes. If this algorithm were seen to have any biases towards or against any specific group of applicants, they would first run the risk of being faced with lawsuits, regulatory penalties, and public backlash, but along with that, they would also have to invest a significant amount of resources to address these issues. The perfect example for this situation is the Amazon hiring bias situation mentioned earlier; Amazon's AI algorithm held biases when it came to male applicants over female applicants, and because this situation came to light, the company faced public backlash and had to scrap the algorithm altogether, resulting in wasted development costs on top of the decline in reputation.

Artificial intelligence has many benefits to a company for increased efficiency, but by no means does it not have any drawbacks. The significant negative effects for artificial intelligence being implemented falls onto both companies and people. Certain groups of people can be discriminated against, perpetuating the historical cycle of human bias.

Conclusion

The introduction of artificial intelligence (AI) into production and the workforce has fundamentally reshaped the job market in numerous different ways. Automation has found itself in many forms of production and businesses throughout all of world history, but AI may be the most severe form of automation yet, bringing along both the good and the bad with it. Overall, the impact of artificial intelligence on anything, especially in the job market, varies and depends on many factors. To understand this deeper, we have to look at both the positives and the negatives of AI implementation.

Looking at the impacts of AI, we can start by looking at what people consider to be the number one issue: AI's effect on unemployment rates. On a broad level, people are more likely to think that the implementation of AI into businesses and firms will only cause jobs to be taken by AI, and therefore, more people will lose their jobs or be less likely to get one in certain industries, but this isn't entirely true. While yes, jobs will be lost to AI, there are also plenty of jobs that now exist and will be open because of AI. Low-skill jobs are more vulnerable to automation because they involve repetitive tasks, while high-skill jobs are less likely to be

replaced by AI (in the near future, at least) because these positions generally require more education and specialized training compared to low-skill jobs. Positions that require people to manage or train AI systems will see an increase in the future. Overall, jobs will be impacted in many ways, making it hard to define a certain trajectory, but most likely, the number of jobs will decrease due to the increased capability of AI, the increased skill mismatches that AI causes, and the ratio between AI managing roles to the jobs that AI takes.

Additionally, bias in AI can have a significant impact on the labor market. The use of AI-based hiring algorithms runs the risk of introducing biases that can affect certain groups of people, giving advantages to some while disadvantaging others. While not all AI have demonstrated significant biases against certain people, some have, and the general risk of bias is an overall negative impact of AI. Some may argue that it's okay to run the risk of AI bias because of the increased efficiency that AI can provide. The main reason businesses and firms use AI-based hiring algorithms is simply because of how efficient they can be, saving money and critical work time. This creates a trade-off between mitigating AI bias and pursuing operational efficiency.

Ultimately, the impact of AI in the workforce remains complex and context-dependent, requiring ongoing evaluation and strategic adaptation. In some ways there are negative impacts, like future unemployment in low-skill jobs and biases in hiring, but there are also positive impacts, such as increased efficiency in production and even creating some jobs for people. The most likely outcome of further implementation will be a decrease in employment for low-skill tasks, and more people will adapt to jobs consisting of high-skill tasks that require more specialized training.

Artificial intelligence presents both opportunities and drawbacks to the labor market and economic inequality. In order for people to adapt to the changes that are being made, it's important to understand the different types of impacts and how AI is inherently different from historical automation trends. This paper identifies persistent bias and the theoretical limits to the capabilities of AI as key factors to understand this impact.

References

M.R. Frank, D. Autor, J.E. Bessen, E. Brynjolfsson, M. Cebrian, D.J. Deming, M. Feldman, M. Groh, J. Lobo, E. Moro, D. Wang, H. Youn, & I. Rahwan, Toward understanding the impact of artificial intelligence on labor, *Proc. Natl. Acad. Sci. U.S.A.* 116 (14) 6531-6539, <https://doi.org/10.1073/pnas.1900949116> (2019).

Grennes, T. (2003). *Creative destruction and globalization*. Columbia International Affairs Online. https://ciaotest.cc.columbia.edu/olj/cato/v22n3/cato_v22n3grt01.html

Koenig, M. L., Oseid, J. A., & Vorenberg, A. (2019). *Ok, google, will artificial intelligence replace human lawyering?*. Marquette Law Scholarly Commons. <https://scholarship.law.marquette.edu/facpub/709/>

Tai, M. C.-T. (2020, April 9). *The impact of artificial intelligence on human society and... : Tzu Chi Medical Journal*. Tzu Chi Medical Journal. https://journals.lww.com/tcmj/fulltext/2020/32040/the_impact_of_artificial_intelligence_on_human.5.aspx

Rhue, L. (2018, December 6). *Racial influence on automated perceptions of emotions*. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3281765

Bessen, J. (2019, June 7). *10. Artificial Intelligence and jobs: The role of demand*. De Gruyter. <https://www.degruyter.com/document/doi/10.7208/9780226613475-012/pdf?licenseType=restricted>

Bughin, J., Hazan, E., Lund, S., Dahlstrom, P., Wiesinger, A., & Subramaniam, A. (2018, May 23). *Skill shift: Automation and the future of the workforce*. McKinsey & Company. <https://www.mckinsey.com/featured-insights/future-of-work/skill-shift-automation-and-the-future-of-the-workforce>

Department of Medical Sociology and Social Work. (n.d.). *The impact of artificial intelligence on human society and... : Tzu Chi Medical Journal*. LWW. https://journals.lww.com/TCMJ/fulltext/2020/32040/The_impact_of_artificial_intelligence_on_human.5.aspx

Lauren Rhue Assistant Professor of Information Systems and Analytics. (2024, June 5). *Emotion-reading tech fails the racial bias test*. The Conversation. <https://theconversation.com/emotion-reading-tech-fails-the-racial-bias-test-108404>

Lambrecht, A., & Tucker, C. (2019, April 10). *Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of STEM career ADS | management science*. INFORMS PubsOnline. <https://pubsonline.informs.org/doi/10.1287/mnsc.2018.3093>



Ukanwa, K., & Rust, R. T. (2018). *Discrimination in Service*. Marketing Science Institute.
http://thearf-org-unified-admin.s3.amazonaws.com/MSI/2020/06/MSI_Report_18-121-1.pdf

Dressel, J., & Farid, H. (2018, January 7). *The accuracy, fairness, and limits of predicting recidivism* | *science advances*. Science Advances.
<https://www.science.org/doi/abs/10.1126/sciadv.aao5580>

Weissmann, J. (2018, October 10). *Amazon created a hiring tool using AI. it immediately started discriminating against women*. Slate Magazine.
<https://slate.com/business/2018/10/amazon-artificial-intelligence-hiring-discrimination-women.html>

Wilson, H. J., & Daugherty, P. R. (2017, March 23). *The jobs that artificial intelligence will create*. MIT Sloan Management Review.
<https://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/>