# THE IMPACTS OF USING AI-BACKED APPLICANT TRACKING SYSTEMS FOR RECRUITMENTS PROCESSES

A Research Paper submitted to the Department of Engineering and Society
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Systems Engineering

By

Vienna Donnelly

March 25, 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.

ADVISOR
Catherine D. Baritaud, Department of Engineering and Society

**THE RISE OF ALGORITHMIC BASED DECSION-MAKING IN A DATA RICH WORLD**

The second half of the 20th century brought about substantial change and innovation to the realm of technology, which, in turn, revolutionized the field of analytics. Corporations and organizations began to actively document and record their business processes in an effort to better comprehend the state of their businesses ushering in a new age of data collection and analysis (Davenport, 2013, p.66). At its inception, data analysis was reserved for organizations that possessed the infrastructure and manpower to collect and meaningfully process large quantities of data. Today, however, Davenport (2013) explains that information firms are not the only organizations capable of harnessing data, “it’s every firm in every industry. If your company makes things, moves things, consumes things, or works with customers, you have increasing amounts of data on those activities” (p. 67). The field of analytics aims to harness vast quantities of data and utilize the results to frame future decisions. This goal is often accomplished by discerning and modelling predictive statistics or by summarizing and contextualizing endless data into a useful set of metrics. In addition to these methods, decision algorithms can be made to classify data and generate automated responses to a pre-determined question. While the benefits of harnessing data to make well-informed decisions are clear, what are the harms? A recent Pew Research study indicated that 58% of Americans believe that algorithms possess the same biases as their creators (Pew Research Center, 2018, p. 8). The presence of bias within data and algorithms has the potential for long lasting ramifications within the organizations using these techniques. For this reason, it is necessary to not only recognize the possibility of bias within machines but also to determine methods of mitigation.

Sports analytics and hiring analytics are two notable applications of the use of data driven decision-making in institutions that previously relied primarily on human decision-making. In 2002 Billy Beane, the general manager of the Oakland Athletics, used data analysis to recruit an unusual lineup that fit his minimal budget and transformed the Oakland A's into a winning team with the longest winning streak in a century (Steinberg, 2015, para. 2). Following Beane's success, the application of analytics erupted in the world of sports and today, "every major professional sports team either has an analytics department or an analytics expert on staff" (Steinberg, 2015, para. 3). Beginning in the Fall semester, the Capstone team of Michael Bassilios, Joshua Barnard, Vienna Donnelly, Rachel Kreitzer, and Ava Jundarian started work on a recommendation system for the creators of GameForge, a university-level golf training program. The technical project analyzed tournament data for junior and college golfers in order to create a set of models that will help coaches gain insight on which junior golfers will be the best fit and have the highest potential for improving their team. The models will culminate in a dashboard designed to indicate preferential players for college teams.

Recommendation algorithms have also become commonplace within the recruiting protocols of numerous human resource departments. Many companies have turned to applicant tracking systems (ATS) to filter applicant pools to a manageable size (Black & van Esch, 2020, p. 216). Loosely coupled with the technical project, the Science, Technology, and Society (STS) research project aims to detail the presence of bias within algorithms, the use of analytics within the hiring space and potential mitigation tactics against bias. By using the framework of Social Construction of Technology (SCOT) to understand the manifestation of bias in a recommender system, the thesis will answer the research question of whether implicit biases can be detected and prevented in recommendation systems (Bijker, Hughes, Pinch, & Douglas, 2012, p.22).

## THE EXISTENCE OF BIAS WITHIN THE FIELD OF ARTIFICIAL INTELLIGENCE

In a 2015 progress report, the White House notes that, “there also exists the potential for big data technology to be used to discriminate against individuals, whether intentionally or inadvertently, potentially enabling discriminating outcomes, reducing opportunities and choices available to them” (Executive Office of the President, 2015, p. 6). This report comes as a governmental response after several documented instances of bias within widely used applications of artificial intelligence, like bias in Google’s search engine technology and bias within several facial recognition technologies (Howard & Borenstein, 2018, p. 1524-1526).

To understand the presence of bias within artificial intelligence (AI), it is critical to understand what biases are and in what forms they can present themselves. Simply put, bias is the predisposition of a person to think about or treat another differently based on their characteristics, specifically in regards to a person’s sex, age, or race (Howard & Borenstein, 2018, p. 1522-1523). While it is possible for outward biases to affect AI, it is much more common for biases in AI to be implicit. For example, a hiring manager could have the opinion that younger employees are more desirable and as a result sets a hiring algorithm to filter out all job applicants over the age of 40. It is far more likely, however, that a manager calibrates an algorithm to prioritize a certain skill that they might unconsciously associate with younger employees, like the ability to write code in a brand-new programming language, thus resulting in an age-biased subset of the original applicant pool.

Decision analytics are often derived from large sets of historical data and as algorithms learn from data, “they find patterns within datasets that reflect our own implicit biases and, in so doing, emphasize and reinforce these biases as global truth” (Howard & Borenstein, 2018, p. 1524). For example, one avenue of incorporating bias into an algorithm stems from biased sampling, “in which groups are over or underrepresented in the training data” (Manyika, Silberg, & Presten, 2019, para. 4). Another inadvertent method of introducing bias into an algorithm can come from using training data that can “ include biased human decisions or reflect historical or social inequities” (Manyika, Silberg, & Presten, 2019, para. 4).

## THE APPLICATION OF ARTIFICIAL INTELLIGENCE IN THE HIRING SPACE

Using artificial intelligence to power applicant tracking systems, companies have been able to reduce both the time and costs associated with recruiting, making the adoption of Applicant Tracking Systems (ATS) essential for a prospering business (Black & van Esch, 2020, p. 218). In 2019, 99% of Fortune 500 companies used an applicant tracking system to manage their talent acquisition, making the use of ATS a prevalent solution for companies searching for human capital (Qu, 2019, para. 5). Concern, however, exists regarding the possibility for unintended bias to exist within applicant tracking systems and subsequently a company’s personnel. Considering the prevalence of applicant tracking systems in today’s hiring processes and their expected market growth despite the relative “newness” of artificial intelligence and machine learning technology, it is feasible to assume that use of applicant tracking systems will continue to grow and impact job acquisition (Markets and Markets, 2018, para. 1). It is clear that ATS are influential in the hiring process. Therefore, any biases present in applicant tracking systems, whether by intentional programming or unintentional, implicit bias, have the potential to critically heighten the presence of bias or discrimination within an applicant pool and, eventually, a company.

The theory of the Social Construction of Technology as described by Bijker et al. (2012) is a multidirectional approach to categorizing the development of a technical artifact through the analysis of the artifact’s relevant social groups (Bijker, Hughes, Pinch, & Douglas, 2012, p.22).

Bijker et al. (2012) describe this development as "an alternation of variation and selection" (Bijker, Hughes, Pinch, & Douglas, 2012, p.22). Using this theory, the ways in which applicant tracking systems are shaped by relevant social groups can be discerned. SCOT can also explain how an applicant tracking system shapes these same groups. As detailed in Figure 1, ATS technologies interact with many social groups including the developers creating the system, the human resources executives commissioning the technology, the applicants to a company, and the managers seeking new personnel. ATS are also influenced by cultural ideas like a company's culture or the requirements and skills for an open position. Developers of commercial applicant tracking systems take in the objectives and goals of a company's executives when creating a system in order to provide the company with a set of applicants that best fit their overall interests. Application tracking systems are often developed using algorithms which assess data on the historical success of a company as well as current company culture in order to predict which persons in an applicant pool will serve as the most

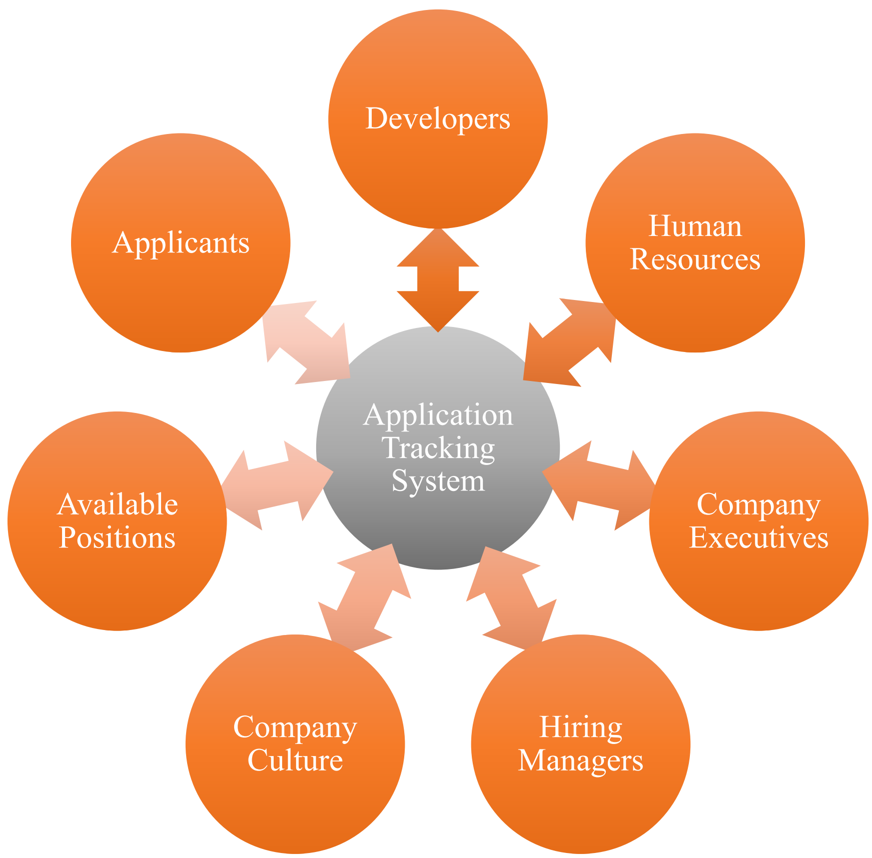


Figure 1: The Social Construction of an Applicant Tracking System: A depiction of the many relevant social groups whose feedback alters the development and use of an application tracking system (Donnelly, V., 2021).

prosperous employees for the company. In turn, the system which alters the set of available applicants affects the future of the company. By shaping the business's future personnel, applicant tracking systems have a tremendous effect on the future of a company.

While it is the goal of a developer to provide a "good" ATS that will compute a desirable set of applicants, there are tradeoffs in algorithms like this. A main tradeoff is the one between certainty of fit and breadth of search. A developer can adjust the output of their decision-making algorithm to generate a limited set of applicants that they can say perfectly meet the objectives and goals of a position with a degree of certainty. In this method a developer will inevitably miss many sufficient candidates but will not recommend any poor candidates. Conversely, a developer or system that values a wide breadth of candidates will ensure it recommends all sufficient candidates in a pool, inevitably recommending a few poor candidates as well but the system will not miss any of the desirable candidates. Tradeoffs between these two types of errors, commonly classified and Type I and Type II errors, are inevitable within artificial intelligence and can magnify the presence of bias within a system.

It is important to note the process these systems follow in order to transform a company's hiring goals into a system that can track and classify applicants in order to identify shortcomings and areas that can be manipulated. Previous iterations of applicant tracking systems simply scanned resumes, looking for keywords that matched the desired skills for a certain position. Current applicant tracking systems, however, contain many levels of filtration. The ATS reads in resume data and separates experiences into different categories, such as experience and education. The system then matches the newly read in data to keywords set by a recruiter for the position in question and finally scores applicants based on their aptitude for the position and company (Savage & Bales, 2017, p. 215-216). An estimated 75% of resumes never make it past the applicant tracking system and into the hands of a recruiter, making the use of ATS a significant barrier to entry for candidates whose resumes are not chosen (Bell, 2018, para.1). The first hurdle for applicants is having their experiences recognized as significant by the system which can be inhibited by unique formatting, wording or font choice. Next is being considered a “good fit” for the role. Many companies do not score candidates solely on the presence of desirable words in their resume but rather by using predictive analytics to classify candidates based on certain attributes which the company will then use to support their decisions in identifying the candidates with the potential to be most successful ( Savage & Bales, 2017, p. 216 - 224).

There does not exist an objective set of traits, experiences, or skills that can indicate or predict superior job performance. Therefore, all algorithms used to identify superior candidates are based on opinion or precedent, and biases, whether outright or implicit, present in the organization can become exacerbated (Howard & Borenstein, 2018, p. 1521). This indicates that the technology itself does not determine who will be considered a desirable candidate, but rather humans designing and influencing the system will. Consequently, the interpretations of this technology are flexible. Companies who are gaining valuable human capitol despite a reduction in workload and costs will have a greater opinion of ATS when compared to applicants who feel that they or their minority groups are not being recognized by the system.

For example, the lack of diversity present within major tech companies in Silicon Valley has long been questioned and is frequently met with the argument that these companies are hiring for “cultural fit” (Barton, 2019, para. 9-13 ). This lack of diversity, however, could be explained by algorithmic bias. Research from a 2017 American Bar Association journal explains, “Algorithm creators rely on employers’ past hiring data to build predictive formulas to match their best workers' traits with job applicants’ data. If a company has not historically hired people from a particular racial or ethnic group, its algorithm will systematically exclude such people from consideration” (Savage & Bales, 2017, p. 224 ).

## DETECTING AND REGULATING BIAS IN APPLICANT TRACKING SYSTEMS

Bias within AI has the potential to threaten the development and growth of many areas within the world including the workforce. Unaddressed, this issue will at best, imitate the current state of the world, and at worst, extract and magnify the implicitly bias choices which exist in everyday life. There have been many instances of bias in artificial intelligence affecting everyday life, for example, a 2018 study conducted by Timnit Gebru and Joy Buolamwini found that certain facial recognition software correctly recognized 99% of white males but only recognized 35% of black females, as such this is an issue that must be addressed (Tiku, 2020, para. 22). Fortunately, the most effective way of mitigating bias in artificial intelligence is to recognize its manifestation in an algorithm and account for it in the decision-making process.

Many researchers in the field of AI believe that a first step in preventing bias is to form multidisciplinary teams which include members with expertise regarding bias in society and machines (Howard & Borenstein, 2018, p. 1532). Howard and Borenstein (2018) recommend that “those trained to recognize and mitigate implicit bias should be an integral part of the design process” (p. 1532). In addition, many methods exist to determine whether or not a system produces biased results. An initial check on the legitimacy of a set of ATS results would be to compare the distributions of an input set and an output set for variables that are commonly seen as discriminating factors like sex, age or race. A statistically significant shift in the distributions of inputs and outputs is an indication that a certain discriminating factor played a role, whether outright or inadvertently, in the decision-making process. Another method of evaluation presented by Howard and Borenstein (2018) in an example regarding bias in medical care involves the evaluation of discriminatory factors as independent variables stating, "Processes and outcomes could be evaluated using patient race or gender as the independent variable to ensure there are no significant differences found in the dependent variables of decisions made or quality of care when changing the value of this parameter" (p. 1532).

## COMBATTING BIAS IN MACHINES

While many methods of detecting bias within artificial intelligences exist, the effort to detect and confront bias within machines before commercialization is not universally made. In 2020 for example, Google recruited Timnit Gebru, a prominent researcher in the field of unethical applications of artificial intelligence, to run a team conducting research on AI (Tiku, 2020, para. 1). In her time at Google, Gebru and her team began researching the potential ethical concerns surrounding large language models, which are significant for the accurate processing of Google search requests. Her team authored a research paper detailing four possible categories of harm associated with large language models, with the expectation of publishing the research. After circulating the draft, Gebru received a call from Google's research Vice President asking for a retraction of the research paper. Ultimately, this series of events resulted in the termination of Gebru's work with Google, an event that many within the company have called an act of retaliation (Tiku, 2020, para. 15). This example of a company's struggle to accept the ethical implications of their technology demonstrates the lack of a key player within the realm of artificial intelligence and thereby the realm of application tracking systems. In order to address the prevalence of bias in AI technologies, the field requires a common set of standards for practice and implementation which can be enforced with the introduction of a regulator to the system.

As is clear in the example of Gebru's departure from Google, a private company is not required to enforce accountability within their developmental practices. The introduction of a third-party regulator, backed by a multidisciplinary group of experts in bias, analytics and AI, could require an analysis of bias in new technologies and confront glitches before products are commercialized and distributed to the public.

The previously introduced model of the development of an application tracking system using the framework of the social construction of technology depicted in Figure 1. on page 6 aptly shows the relevant social groups and pertinent ideals which play a role in the construction of an ATS. This framework, however, does not sufficiently describe the hierarchy between groups present in this system. The ANT Handoff model presented by Carlson (2007) can more linearly depict each of the groups needed to transform an idea into a product ready for use (Carlson, 2007, p. 2). This linear depiction can be used to show the definitive effect a regulator might have on the overall systems. Figure 2 describes the current interactions between relevant groups within the system. This model describes the flow of work from party to party, starting

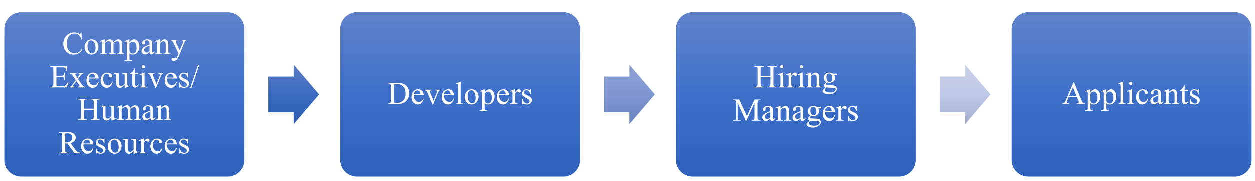


Figure 2: The Handoff Model for the Development and Use of an Applicant Tracking System: A depiction of the four actors who currently play roles in the development and results of an application tracking system (Donnelly, V., 2021).

with a company's executives commissioning an applicant tracking system by developers who then provide a recommendation system for hiring managers. Managers decide when to use ATS for new positions, opening it to new applicants, and then implement the decision-making capabilities of the technology to make hiring decisions. At present, there are no parties working on behalf of the applicants. The addition of a regulator who would assess the validity of an algorithm's decisions would act as a protective party of the applicants who are one of the end users of this technology. Figure 3 demonstrates the effect of introducing a regulator to the process flow in the development and use of an applicant tracking system. This additional actor,


Figure 3: The Modified Handoff Model for the Development and Use of an Applicant Tracking System: A depiction of the addition of a regulatory actor in the system and the effects the additional actor has in the development and results of an application tracking system (Donnelly, V., 2021).

indicated by a red outline, has the duty and ability to identify and prevent implicit bias from affecting the results of the decision-making algorithm powering an applicant tracking system.

A case study detailed by Sammy (2019) illustrates the reality of incorporating a regulatory actor into the development of a decision-making algorithm. A credit card company wanting to determine a subset of their customers to offer a preferred rate using a classification algorithm. The algorithm learned its decision rules from a set of data containing equal proportions of men and women, however, the auditor reviewing the algorithm noted that only 20% of women were being selected to receive a preferred rate compared to 50% of men. Furthermore, the auditor reported that only 6% of customers over 50 were offered the preferred rate. After removing inappropriate classification variables at the behest of the auditor, the credit card company implemented an adjusted classification algorithm which predicted 40% of women and 50% of men as good candidates for the preferred rate. The adjusted algorithm also improved its classification of individuals over 50 to 21% (Sammy, 2019, p.45-46). Through the use of a third-party regulator, the credit card company in this example was able to considerably lessen the effects of bias within their classification algorithm while still obtaining meaningful results for the selection of preferred rate customers.

**PURSUING AN UNBIASED FUTURE**

Technological advancements in the practice of using data analytics coupled with artificial intelligence algorithms as a decision-making tool have greatly impacted the areas in which these techniques are implemented. Positive impacts of AI implementation include reductions in time and manpower needed to complete certain tasks, expansions in the understanding of complex systems, and developments in predictive technology. Reports of instances of bias in algorithms used in facial recognition technology, search engines, healthcare technology, and many other areas have spurred multiple investigations into the ethics of artificial intelligence. Artificial intelligence has many applications in the modern world, as such, the existence of bias within these technologies can have widespread consequences. As a result, it is essential to recognize the prospect of bias within artificial intelligence and employ safeguards to combat it. The addition of a regulatory actor, equipped to recognize bias and offer specific solutions, in the process of developing decision-making algorithms is a crucial step towards producing unbiased technology. Failure to act on this issue can lead to the unintentional permeation of bias in systems using algorithmic decision-making tools. By making the effort to identify and resolve potential occurrences of bias, however, developers of AI algorithms can have a positive impact on both the social and technical environments of the spaces in which they are implemented.

**REFERENCES**


Barton, L. (2019, March 1). Silicon Valley: Using culture fit to disguise discrimination? *Raconteur.* Retrieved from https://www.raconteur.net

Bijker, W. E., Hughes, T. P., Pinch, T., & Douglas, D. G. (2012). *The social construction of technological systems: New directions in the sociology and history of technology*. Cambridge, Ma: MIT press.

Black, J. S., & van Esch, P. (2020). AI-enabled recruiting: What is it and how should a manager use it? *Business Horizons*, *63*(2), 215–226. doi:10.1016/j.bushor.2019.12.001

Bell, T. (2018, April 17) The secrets to beating an applicant tracking system (ATS). *CIO*. Retrieved from https://www.cio.com

Carlson, W.B. (2007). STS research models. Adapted by C. Baritaud (2009) for *STS 4500/4600 thesis modules*. University of Virginia, Charlottesville, VA.

Davenport, T. H. (2013, December). Analytics 3.0. *Harvard Business Review,* 91(12), 64-72. Retrieved from https://hbr.org

Donnelly, V. (2021). *The social construction of an applicant tracking system*. [Figure 1]. *STS Research Paper: The impacts of using AI-backed applicant tracking systems for recruitment processes* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Donnelly, V. (2021). *The handoff model for the development and use of an applicant tracking system*. [Figure 2]. *STS Research Paper: The impacts of using AI-backed applicant tracking systems for recruitment processes* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Donnelly, V. (2021). *The modified handoff model for the development and use of an applicant tracking system*. [Figure 3]. *STS Research Paper: The impacts of using AI-backed applicant tracking systems for recruitment processes* (Unpublished undergraduate thesis). School of Engineering and Applied Science, University of Virginia. Charlottesville, VA.

Executive Office of the President. (2014, May). *Big data: Seizing opportunities, preserving values*. Retrieved from https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf

Howard, A., & Borenstein, J. (2018). The ugly truth about ourselves and our robot creations: The problem of bias and social inequity. *Science and Engineering Ethics*, *24*(5), 1521–1536. doi: 10.1007/s11948-017-9975-2

Manyika, J., Silberg, J., & Presten, B. (2019, October 25). What do we do about the biases in AI? *Harvard Business Review*. Retrieved from https://hbr.org

Markets and Markets. (2018, July) *Applicant tracking system market size, share and global market forecast to 2023*. Retrieved from https://www.marketsandmarkets.com/Market-Reports/applicant-tracking-system-market-27004100.html

Pew Research Center (2018, November 16). *Attitudes toward algorithmic decision-making*. Retrieved from https://www.pewresearch.org/internet/2018/11/16/attitudes-toward-algorithmic-decision-making/

Qu, L. (2019, November 7). 99% of Fortune 500 companies use applicant tracking systems (ATS) [Blog post]. Retrieved from https://www.jobscan.co/blog/99-percent-fortune-500-ats/

Sammy, A. (2019). Bias in the machine. *Internal Auditor*, *76*(3), 42-46.

Savage, D. D., & Bales, R. (2017). Video games in job interviews: Using algorithms to minimize discrimination and unconscious bias. *ABA Journal of Labor & Employment Law*, *32*(2), 211-228.

Steinberg, L. (2015, August 18). Changing the game: The rise of sports analytics. *Forbes.* Retrieved from https://www.forbes.com

Tiku, N. (2020, December 23). Google hired Timnit Gebru to be an outspoken critic of unethical AI. Then she was fired for it. *Washington Post*. Retrieved from https://www.washingtonpost.com