

# Algorithmic Bias: A New Legal Frontier

**Ken Serwin**  
*Berkeley Research Group, LLC*  
Emeryville, California

**Alfred H. Perkins, Jr.**  
*Starnes Davis Florie LLP*  
Birmingham, Alabama

Reliance on algorithms to process data and make decisions based on that data is increasing. So too are concerns about discriminatory or at least unfair decisions as a result of algorithms. Decisions perceived to be biased based on race, gender, age or other protected characteristics can lead to and have led to litigation.

We are increasingly living in a Matrix that most of us do not perceive.<sup>1</sup> Chander notes a line of recent scholarship attacking the increasing role of automated algorithms in our lives. Some are concerned that the rise of algorithmic decision-making will make discrimination even more difficult to ferret out, while others argue that algorithms *reduce* the invidious discriminations that result from human decision-makers with unfettered discretion.<sup>2</sup>

## What is an Algorithm

An algorithm is a specific set of instructions or finite sequence of steps used to solve a problem or accomplish a task.<sup>3</sup> In today's society, algorithms have been integrated into decision making processes for a wide range of institutions.<sup>4</sup> In general, there are two types of algorithms: computational and learning.

## Computational vs Learning Algorithms

The first type of algorithms, computational algorithms, simply apply the previous definition to computing systems. This type of algorithm essentially applies a set of steps or rules, based on some logic, to an input to yield an output. In the past, computational algorithms were used, for example, to power trading models as detailed, instruction-based models that analyzed "clearly defined data and variables."<sup>5</sup> Computational algorithms are commonly used in computer programs as a means to analyze sets of data and return solutions to a specific question.<sup>6</sup>

The second type of algorithms, learning algorithms, have been developed to advance computational algorithms for the purpose of artificial intelligence.<sup>7</sup> The goal of learning algorithms is to not only examine inputs, but to further learn from the information that is examined and use what has been learned to evolve how the system interprets its inputs.<sup>8</sup> While a learning algorithm contains sequential computational

---

<sup>1</sup> Anupam Chander, *The Racist Algorithm?*, 115 MICH. L. REV. 1023, 1024 (2017)

<sup>2</sup> Anupam Chander, *The Racist Algorithm?*, 115 MICH. L. REV. 1023, 1027 (2017)

<sup>3</sup> Osoba, Osonde and Welser IV, William, "An Intelligence in Our Image," Rand Corporation, 2017, pp. 4-5; <https://www.merriam-webster.com/dictionary/algorithm>

<sup>4</sup> Petrasic, Kevin, et. Al. "Algorithms and bias: What lenders need to know," White& Case LLP, 2017, p. 1.

<sup>5</sup> Petrasic, Kevin, et. Al. "Algorithms and bias: What lenders need to know," White& Case LLP, 2017, p. 2.

<sup>6</sup> Osoba, Osonde and Welser IV, William, "An Intelligence in Our Image," Rand Corporation, 2017, pp. 4-5.

<sup>7</sup> Petrasic, Kevin, et. Al. "Algorithms and bias: What lenders need to know," White& Case LLP, 2017, p. 2.

<sup>8</sup> *Ibid.*

algorithms, it can adjust its behavior based on its “experience,” which it gathers from inputs.<sup>9</sup> For instance, a university may use a learning algorithm to compare applications to performance as a means of determining which applicants have the best chance of succeeding at the school and output the top applicants for in-person interviews.

## Concerns about Algorithmic “Bias”

As the use of algorithms, particularly learning algorithms, has become more pervasive and integral to the decision making processes within institutions, concerns about fairness and bias in algorithms have surfaced.<sup>10</sup> Algorithms used to support financial scoring, hiring, and sentencing decisions are just a few examples of uses that have raised concerns and questions about fairness and lawfulness.

Consider financial scoring algorithms, for example. Such algorithms provide financial institutions with a “score” that represents the customer or applicant’s suitability for a financial product such as a loan.<sup>11</sup> In the past, more traditional algorithms were used to assess an individual’s credit score based on measures of repayment such as debt-to-income or loan-to-value ratios.<sup>12</sup> Today, learning algorithms are frequently used for this purpose. However, today’s algorithms may include less traditional inputs (i.e. search histories or shopping patterns) in order to analyze the creditworthiness of potential borrowers with limited access to traditional input values.<sup>13</sup> Some of the additional inputs considered may contain underlying biases that could affect lending decisions. Furthermore, accuracy of this additional data may be more suspect, as it may not be provided directly by the potential borrower. As a result, these scoring algorithms may be susceptible to outcomes that reflect discrimination and lending biases if they are not monitored correctly.<sup>14</sup>

Similarly, employers’ use of hiring algorithms could potentially insert bias into the hiring process. Many companies use algorithms to sort through applicants and determine the best candidates for a specific job. Some companies even outsource this procedure directly to third parties. However, algorithms that help with hiring decisions could inadvertently rely on data samples or decision rules that are not reflective or representative of the actual universe of employees. For example, if a hiring algorithm based selection on a learned decision factor such as attendance at highly selective universities, then the output of such an algorithm may be susceptible to any underlying biases in the selection criteria of those universities. If, for example, the universities were to only take students from certain neighborhoods, then the hiring algorithm could be biased against the neighborhoods from which the university did not include.<sup>15</sup>

Sentencing algorithms, used in the criminal justice system to provide risk assessments and inform sentencing decisions are another instance in which the lawfulness and accuracy of algorithms has been questioned.<sup>16</sup> Angwin et. al. (2016) discusses a study of the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) system software used by Northpointe to assess criminal risk in

---

<sup>9</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, pp. 4-6.

<sup>10</sup> “Taking Algorithms to Court,” AI Now Institute, September 24, 2018: “Litigating Algorithms: Challenging Government Uses of Algorithmic Decision Systems,” AI Now Institute, September 2018, p. 1; MacCarthy, Mark, “Standards of Fairness for Disparate Impact Assessment of Big Data Algorithms,” Cumberland Law Review, 48 – 102, April 2, 2018. Available at SSRN: <https://ssrn.com/abstract=3154788>

<sup>11</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 3.

<sup>12</sup> Ibid.

<sup>13</sup> Ibid.

<sup>14</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 3.

<sup>15</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 4-

5.

<sup>16</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, p. 13.

sentencing and parole decisions. The results of this study pointed to systematic racial bias when determining the recidivism risks of convicts.<sup>17</sup> In this circumstance, the COMPAS system was adapting due to specific learning procedures within the algorithm, and as a result, the algorithms output suggested a bias in favor of non-black convicts even when their offenses were more severe than black convicts.<sup>18</sup>

Potentially discriminatory outcomes described in the above instances are examples of disparate impact, a legal term used to describe the “systematic disadvantages artificial agents impose on subgroups based on patterns learned via procedures that appear reasonable and nondiscriminatory on face value.”<sup>19</sup> Disparate impact is one of the main algorithmic outcomes that has led to questions of bias regarding the use of algorithms for decision-making processes.

## Causes of Algorithmic Biases

There are three main causes of algorithmic bias: input bias, training bias, and programming bias.<sup>20</sup> On the other hand, algorithmic outcomes oftentimes labeled as “biased” may simply reflect unpleasant facts based on causal relationships derived from reliable representative data.

Input bias can be produced when inputs are non-representative, lack information, historically biased, or otherwise “bad” data.<sup>21</sup> Algorithms are subject to the inputs they analyze and as a result, biases inherent to inputs are often preserved by an algorithm, and even propagated in the outputs. In this way, “an artificial agent [or algorithm] is only as good as the data it learns from,”<sup>22</sup> or more descriptively – crap in, crap out.

Training bias is similarly linked to the data used in an algorithm. Training bias can arise where input data is mis-categorized or where outputs are inappropriately assessed.<sup>23</sup> Often, data is used to “train” an algorithm so that it can learn from its own experience. However, learning algorithms can learn from correlation as opposed to causation because they cannot distinguish between the two and do not understand when additional data may be required to produce accurate results.<sup>24</sup> When an algorithm categorizes data and results based on correlations it finds in the data, it can “ignore or obscure other important (and more relevant) factors.” In this way, training bias can be seen when the output of an algorithm is based on certain learned correlations while a different, and potentially more accurate output, may have been produced had the algorithm considered different or additional information.

Programming bias, describes biases engrained in the design of the algorithm such as subjective rules programmed into an algorithm.<sup>25</sup> Additionally, in the case of learning algorithms, programming biases can develop from an algorithm’s interactions with human users and assimilating to existing or new data.<sup>26</sup> For example, the hiring algorithm described above was programmed to favor applicants from specific colleges and was thus biased toward applicants that did not attend the favored universities. A learning algorithm could similarly learn this behavior through human interaction if the users of the algorithm were continuously choosing applicants from specific universities and these choices were an input in the

---

<sup>17</sup> Ibid.

<sup>18</sup> Ibid.

<sup>19</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, p. 11.

<sup>20</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 2.

<sup>21</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 1

<sup>22</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, pp. 4-5; <https://www.merriam-webster.com/dictionary/algorithm>

<sup>23</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 2.

<sup>24</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 4.

<sup>25</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 2.

<sup>26</sup> Petrasic, Kevin, et. Al. “Algorithms and bias: What lenders need to know,” White& Case LLP, 2017, p. 2.

algorithm's data. In this case, the algorithm would begin to understand that the human prefers applicants from a specific school and perpetuate this bias by favoring applicants from that school.

What these causes of algorithmic bias have in common is that they are all, in some way, dependent on either biases in the underlying data or in the programming decisions used to analyze that data. In such circumstances, algorithmic biases are not originated by the algorithm, however they are the result of data or programming that includes or reflects biases.

In the above circumstances, "bias" describes a negative outcome where the algorithm is either "wrong" or otherwise "bad." However, other outcomes referred to as algorithmic "bias" are not objectively either wrong or bad. In the case of the hiring algorithm, there may be nothing fundamentally wrong with favoring applicants that attended specific universities. In fact, this practice is common in many businesses. It is understandable that if a company has success with graduates from a certain university, the business may be more inclined to select other applicants from that same school. It also may be that one university prepares its students better for a specific job than another university. In this case, while there is an observable differential in the hiring suggestions of the algorithm, that differential could be a profit maximizing outcome based on real observed causal relationships. Such outcomes are not objectively bad, wrong, or biased. In some instances, then, observed outcomes referred to as algorithmic bias simply reflect unpleasant facts based on causal relationships derived from reliable representative data. In these instances, while these algorithmic results may reflect disparate impact, the results are not objectively wrong. In such cases, the results are not necessarily a cause for concern and the pejorative description of these results as "biased" is more of a subjective determination.

### Cleaning up Algorithmic Bias

Algorithmic bias is a genuine and growing concern as the use of algorithms becomes more widespread. Different approaches have been suggested as a means to address algorithmic bias and correspondingly avoid discriminatory outcomes and other unfair results. These approaches include: "debiasing" data, constraining algorithmic optimization, implementing causation tests, and performing algorithmic audits.

The first approach, "debiasing" data, is a way to avoid algorithmic bias that appears due to unwanted correlations in data.<sup>27</sup> "Debiasing" data uses principal component analysis to remove potential associations in the data so that previous correlations no longer exist.<sup>28</sup> Once the data is free of its previously biased associations, the algorithm can process the uncorrelated data and avoid the previous bias in the results. Debiasing "fixes" the data used by an algorithm. Whether or not the data needs to be fixed depends on whether it accurately reflects the universe of which it purports to be representative.

Consider the sentencing example discussed above. Suppose that recidivism is a direct function of the type of offense committed and that those convicted of violent crimes commit additional offenses more often than those who were convicted of non-violent offenses. Further suppose the data an algorithm uses to estimate the likelihood of recidivism includes examples from one racial group that committed only non-violent offenses while the data for a different racial group included only those convicted of violent offenses. If such a hypothetical algorithm incorporated race into its decision rules, it would most likely return a result that says one race is more likely than the other to commit another offense after release. This answer would be wrong because of the bias in the data. Such data needs to be fixed. If, on the other hand, the data used

---

<sup>27</sup> Wang, Eric, "What does it really mean for an algorithm to be biased?", The Gradient, May 1, 2018, pp. 4-6.

<sup>28</sup> Wang, Eric, "What does it really mean for an algorithm to be biased?", The Gradient, May 1, 2018, pp. 4-6.

by such an algorithm reflects the actual mix of violent and non-violent offenders for all respective races as exists in the population, then such a fix may not necessary from an objective point of view, though it could still predict differential risks of recidivism for other reasons.

Constrained algorithmic optimization is an approach that “[maximizes] immediate utility, under some formal fairness constraint.”<sup>29</sup> When an algorithm is implemented in place of human decision making, constrained optimization is a way to prioritize long-term fairness in the results so that biases within the data are not perpetuated.<sup>30</sup> By choosing a “decision rule” to constrain the algorithm’s perceived unfairness, the algorithm can produce results that balance utility and fairness, and limit biases and disparate impact.<sup>31</sup> However, the cost of limiting the bias in the algorithm is that it may reduce the effectiveness of the algorithm’s outputs.<sup>32</sup> As with any data model, the programming can force adjust the model to “fix” the result in any way that is desired. If the data and the model are accurate, however, such a fix based on subjective desires to achieve a given outcome may be looking to achieve some social or political objective rather than correcting any inherent “bias.”

Another approach, to address perceived algorithmic bias is to perform causation tests within algorithms. Adding the ability to perform causal reasoning would be a significant step in accepting decisions based on algorithms “because automated causal reasoning systems can present clear causal narratives for judging the quality of an algorithmic decision process.”<sup>33</sup> From an appropriate scientific point of view, this fix is essential, and any model that uses decision rules that are not based on causal relationships is likely inappropriate.

Algorithmic auditing is another approach to preventing algorithmic bias. The auditing approach is based on “comparing algorithmic output with expected equitable behavior.”<sup>34</sup> Multiple methods for auditing algorithms have been proposed, including approaches that require “statistical independence between outcomes and protected variables” and tests to confirm that algorithms don’t violate disparate impact laws.<sup>35</sup> Auditing outcomes provides an opportunity to identify whether or not an algorithm will violate relevant concerns and, to the extent necessary, adjust the algorithm in some of the other ways discussed.

### Litigation related to Algorithmic Bias

Like many other issues on which there is a difference of opinion in this country, the use of algorithms to make decisions has found its way into litigation. For example, in *Sandvig v. Sessions*<sup>36</sup>, researchers planned to engage in audit testing of internet real estate, hiring and other websites through the use of bots and fictitious user profiles, and they planned to make their findings public. Although plaintiffs have what some might consider noble intentions in trying to determine where the websites engage in discrimination, plaintiffs’ actions would violate certain terms of service of the targeted websites, and their

---

<sup>29</sup> Wang, Eric, “What does it really mean for an algorithm to be biased?”, The Gradient, May 1, 2018, pp. 16-17.

<sup>30</sup> Ibid.

<sup>31</sup> Wang, Eric, “What does it really mean for an algorithm to be biased?”, The Gradient, May 1, 2018, pp. 16-17.

<sup>32</sup> Wang, Eric, “What does it really mean for an algorithm to be biased?”, The Gradient, May 1, 2018, pp. 17-18.

<sup>33</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, pp. 22.

<sup>34</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, pp. 21-22.

<sup>35</sup> Osoba, Osonde and Welser IV, William, “An Intelligence in Our Image,” Rand Corporation, 2017, pp. 21-22.

<sup>36</sup> 315 F.Supp. 3d 1 (D. D.C. 2018)

conduct allegedly violates federal law; specifically, the Computer Fraud and Abuse Act (CFAA), a law dedicated to deterring the criminal element from abusing computer technology. Therefore, plaintiffs brought suit against Attorney General Jeff Sessions challenging the Access Provision of the CFAA. The Access Provision states that “[w]hoever . . . intentionally accesses a computer without authorization or exceeds authorized access, and thereby obtains . . . information from any protected computer . . . shall be punished as provided in subsection (c) of this section.” 18 U.S.C. § 1030(a)(2)(C). The CFAA provides for a fine and/or imprisonment upon a first violation of the Access Provision.

The plaintiffs in the case are four college professors and a media organization. Plaintiffs are conducting studies to respond to new trends in real estate, finance and employment transactions which increasingly have been initiated on the Internet. Data brokers assemble consumers’ information from a myriad of sources and place consumers into models that include racial, ethnic, socio economic, gender, and religious inferences about them. Plaintiffs asserted that given the history of racial discrimination in housing and employment, this technology may be harnessed for discriminatory purposes. Specifically, they alleged that when algorithms automate decisions, there is a very real risk that those decisions will unintentionally have a prohibited discriminatory effect.

The Court acknowledged that some argue that one way to determine whether algorithms discriminate against members of a protected class is to engage in Outcomes-Based Audit Testing. That testing commonly involves accessing a website or other network service repeatedly, generally by creating false or artificial user profiles, to see how websites respond to users who display characteristics attributed to certain races, genders or other protected classes. Plaintiffs plan to engage or are engaging in such audit testing. For example, two of the professors are investigating whether computer programs that decide what to display on real estate websites discriminate against users based on race or other factors. They will use an automatic data recording technique known as scraping to record the properties that each bot sees on the real estate sites. They will then examine that data to determine whether race-associated behaviors caused the sock puppets to see different sets of properties.

Likewise, two other professors plan to conduct a study to see whether hiring websites’ algorithms end up discriminating against job seekers based on protected statuses like race or gender. They will also use bots to crawl the profiles of a random selection of job-seekers to obtain baseline demographic data, then create fake employer profiles so that they can search for candidates and record how the algorithms rate those candidates. They will ultimately create fictitious job-seeker profiles and have the fictitious seekers apply for fictitious jobs to examine how the algorithms rank the candidates. The media company and its journalists seek to investigate online companies, websites and platforms, including by examining any discriminatory effects of their use of algorithms.

All plaintiffs are aware that their activities will violate the terms of service of certain websites. All intend to use scraping to record data, which is banned by many of the websites the plaintiffs seek to study. Many of the housing websites prohibit the use of bots. All of the hiring websites targeted by plaintiffs prohibit the use of the techniques plaintiffs will employ. Additionally, some of the websites include non-disparagements clauses in their terms of services.

Plaintiffs claim that their investigations constitute protected speech or expressive activity. They also claim that their investigations expose them to the risk of prosecution under the Access Provision of the CFAA. Plaintiffs therefore filed suit against the Attorney General challenging the act on several grounds.

The government moved to dismiss the complaint on several grounds. The district court denied the government’s motion to dismiss and allowed the case to proceed.

In *Chacko v. Connecticut*<sup>37</sup>, Plaintiff asserted race/national origin discrimination among other Title VII claims against the Connecticut Department of Mental Health and Addiction Services, her employer. Chacko is an Asian female of Indian descent. She was employed as a Principal Physician at the Connecticut Valley Hospital. The hospital is for inpatients who are wards of the state and who have psychiatric disorders, substance abuse disorders, or co-occurring disorders.

At all relevant times, three other doctors at the hospital held the position of Principal Physician, two of whom were Asian, and the other Caucasian. The assignments for physicians were made by a trio of administrators, who claim that they made duty assignments using an algorithm that factored in patient acuity, the complexity of the management of the patient's condition, and the turnover of patients. Chacko disputed that the algorithm was ever used.

At the core of Chacko's complaint was her claim that she and other Asian doctors were assigned heavier workloads than their white counterparts. The parties disputed the relative difficulty of work in the three divisions of the hospital. In July 2005, workloads of the physicians were redistributed because a part-time physician was leaving the hospital. Chacko was assigned an additional unit in the redistribution, bringing her workload to a total of four units. Two other doctors, one white on Asian, also carried four units. Chacko contends that the reorganization resulted in increased workloads for Asian physicians, but not for Caucasian physicians, noting that the workload for one Caucasian doctor remained the same while the workload for the other Caucasian doctor decreased. Although Chacko and three other doctors all had four units, Chacko asserted that the four units contained different amounts of work.

The defendant moved for summary judgment claiming that the allegations of heavy workload do not constitute an adverse employment action under Title VII and that, even if it did, the defendant has offered a legitimate, non-discriminatory reason for the assignment of workloads – the algorithm – which Chacko cannot rebut.

The Court determined that Chacko put forth sufficient evidence to satisfy the minimal burden of demonstrating adverse employment action for the purpose of a prima facie case. The defendant advanced as its legitimate, non-discriminatory reason for the work assignments, the fact that the duty assignments were made using an algorithm that factored in patient acuity, the complexity of the management of the patient's condition, and the turnover of patients. Thus, the Court required Chacko to demonstrate that a material issue of fact existed over whether the state's legitimate reason is a pretext for discrimination.

The Court found that Chacko demonstrated that material issues of fact existed regarding pretext. Although the defendant noted that the hospital used an algorithm to assign workload, no evidence of the application of this algorithm had been provided to the Court. Moreover, it did not appear that the algorithm accounted for the apparent disparity in assigning physician assistants. The Court ultimately denied the summary judgment motion as it related to the claim of discriminatory work load assignments.

This case seems to suggest that if an employer is going to defend a decision that was based on an algorithm, the employer needs to provide some description of the algorithm and its decision-making criteria.

In *Munoz v. Orr*,<sup>38</sup> civilian employees brought a Title VII Class Action against the Department of the Air Force, alleging that the employee promotion system used at the Air Force base operated to

---

<sup>37</sup> *Chacko v. Connecticut*, 2010 WL 1330861, No. 3:07-cv-1120 (D. Conn. Mar. 30, 2010)

<sup>38</sup> *Munoz v. Orr*, 200 F.3d 291 (5<sup>th</sup> Cir. 2000)

discriminate against Hispanic males. Jesus Munoz and Manuel Munoz are Hispanic males, brothers, and employed as part of the civilian work force at Kelly Air Force Base near San Antonio, Texas. They brought suit on behalf of themselves and all Hispanic male civilian employees at Kelly, alleging that the promotion system used by the Air Force for civilian employees has a disparate impact on Hispanic males, i.e., that the system results in fewer Hispanic males receiving promotions than would be expected based on the proportion of the civilian work force at Kelly that they comprised.

Civilian employment at Kelly is organized on the General Service Scale (“GS”), a salary and promotion grid in common use throughout civilian federal employment. Each GS level, or grade, represents a salary range. The GS level to which an employee is assigned depends upon such factors as education level, skill level, time in service, and degree of authority of the position he occupies. Over the course of a career in federal civilian employment, an employee may occupy several different GS levels or steps within a GS level. Certain GS levels are not open to employees without particular qualifications (e.g., a college degree or its equivalent). As a general matter, each job opening is allocated to a particular GS level or range of levels, thus setting the maximum salary that position could accrue. A federal civilian job also has skills requirements and responsibilities attached to it that in part define its GS range.

At Kelly, civilian promotions are handled in part by a Merit Promotion Plan that includes an automated system called the “Personnel Placement and Referral System,” or PPRS. Under PPRS, employees need not submit applications for promotions. Rather, as a position becomes available, PPRS considers all eligible employees within the defined area of consideration for the position. PPRS recursively eliminates employees under increasingly specific job requirements until the desired number of candidates is reached. PPRS thus works like a funnel, at first considering all nominally eligible employees for a promotion and then narrowing the field based on successively more detailed requirements until a short, ranked list is generated. Each stage of this narrowing is known as a “Progression Level Factor” or PLF. Ties between employees are broken by reference to appraisal scores, awards, and service computation date, in order. The list is hand-checked and then forwarded to the selecting official, who chooses one of the employees for the promotion.

The automated program was not without subjective elements. At the beginning of the promotion process, three-person teams establish and rank the job skills relevant to the position. This ranked list is called a “Promotion Evaluation Pattern,” or PEP. The PLFs used by the automated program to narrow the field for a given promotion are derived from these PEPs. Furthermore, within the automated PPRS program, ties between eligible employees are broken in part by appraisal scores and awards and service computation dates. An employee’s appraisal scores and receipt of any awards depend, to a large degree, on the discretion of his supervisors. Lastly, after a finite list of names for a promotion has been prepared by the PPRS, a selecting officer chooses one employee from the group. Though the officer’s range of choice is limited to the list derived from the PPRS, the actual selection from within the group is left to the selecting officer’s discretion. Thus, promotions at Kelly comprise both subjective and objective components that are significantly intertwined.

During discovery, plaintiffs sought certain information regarding the Air Force’s promotion procedure, including access to the algorithm used in the automated PPRS process. After an in-camera review of the algorithm, the District Court denied plaintiffs’ request. No objections were filed to that denial. The District Court granted summary judgment against the plaintiffs.

One of the things the 5<sup>th</sup> Circuit reviewed was the District Court’s decision to deny plaintiffs access to the Air Force algorithm used in the PPRS process. The District Court ordered the algorithm sealed following an in-camera review of the computer program in which it determined that the algorithm did not



contain any evidence of discrimination on the part of the Air Force. Plaintiffs did not object at the time the algorithm was sealed, and it was eventually returned to the Air Force. Due to the lack of contemporaneous objection, the 5<sup>th</sup> Circuit only reviewed the District Court’s discovery decision for plain error. Under that standard, the 5<sup>th</sup> Circuit found that the District Court was within its discretion in refusing access to the algorithm. The Court also concluded that it was unlikely that the denial of access to the algorithm unduly prejudiced plaintiffs’ claims. Defendant had already supplied detailed information on the overall promotion system and the inputs used by the PPRS automated system. Thus, denial of access to the algorithm arguably could make it more difficult to identify with specificity the aspects of Kelly’s promotion system responsible for any observed disparate impact. However, the 5<sup>th</sup> Circuit concluded that plaintiffs’ claim did not fail on those grounds and, therefore, a remand to the District Court solely because of its denial of access to the algorithm would be inappropriate.

In 2017, homeowners in Illinois filed suit against Zillow, a real estate data website, alleging that Zillow’s website affected their ability to sell their homes.<sup>39</sup> Plaintiffs contended that Zillow’s website uses “Zestimates” to draw consumers and potential home sellers to Zillow.com. Zestimates are a valuation tool reflecting Zillow’s estimated home value for individual homes determined by a computer algorithm. The Zillow webpage information attached to the First Amended Complaint indicates that a Zestimate “is calculated from public and user-submitted data,” is “a starting point in determining a home’s value” and “is not an appraisal.” The webpage encourages buyers, sellers, and homeowners to supplement Zillow’s information by doing other research such as getting a comparative market analysis (CMA) from a real estate agent and getting an appraisal from a professional appraiser. Plaintiffs assert that Zillow does not inform the general consumer public that it uses Zestimate as a marketing tool to draw home sellers and buyers to its website in an effort to connect them with Zillow’s premier agents. Similarly, Plaintiffs also claimed that Zillow does not tell the general public that it has a financial relationship with the premier agents, including that it makes revenue from advertisements purchased by premier agents and premier lenders displayed on Zillow’s website.

Plaintiffs asserted claims of violation of the Illinois Deceptive Trade Practices Act and the Illinois Consumer Fraud Act. Zillow moved to dismiss. The district court granted Zillow’s motion and dismissed the case. Although Plaintiffs did not prevail in this action, suits like this represent one way in which an individual can dispute the outcomes of an algorithm.

## Other arenas in which Algorithms may affect legal practice.

In addition to spurring litigation, the use of algorithms may affect lawyers in others ways. A recent ABA Journal article reported on the results of pitting lawyers versus an algorithm in spotting issues in non-disclosure agreements.<sup>40</sup> Tasked with spotting issues in five real NDAs from companies including Cargill and Pacific Gas & Electric Co., the software from LawGeex outperformed the attorneys with an average of 94 percent accuracy, according to their study. The attorneys had an average accuracy rate of 85 percent. The attorneys assessed all five NDAs in an average of 92 minutes, while the software took only 26 seconds.

---

<sup>39</sup> *Patel v. Zillow, Inc.*, 2018 WL 2096453, No. 17 C 4008 (N.D. Ill. May 7, 2018).

<sup>40</sup> Jason Tashea, *AI software is more accurate, faster than attorneys when assessing NDAs*, ABA Journal, Feb. 26, 2018, [http://www.abajournal.com/news/article/ai\\_software\\_is\\_more\\_accurate\\_faster\\_than\\_attorneys\\_when\\_assessing\\_ndas](http://www.abajournal.com/news/article/ai_software_is_more_accurate_faster_than_attorneys_when_assessing_ndas).

Algorithms are also being used in the criminal justice system. The use of risk assessment software, powered by sometimes proprietary algorithms, to predict whether individual criminals are likely candidates for recidivism is becoming more common.<sup>41</sup>

In the not-so-distant future, algorithms may help civil-rights agencies predict who could face workplace discrimination. The Equal Employment Opportunity Commission (EEOC) has established the Office of Enterprise Data and Analytics (OEDA). Samuel Christopher Haffer, the first chief data officer hired by the EEOC, said his goal is to create a “distant early warning system,” by, for example, examining groups of similar qualities like race, gender, or ethnicity.<sup>42</sup> The EEOC will employ machine learning- the algorithmic study of past experiences to optimize, or predict, future experiences. The EEOC hopes to be able to flag a particular group of people in a specific industry who would be susceptible to discrimination. The agency would then be able to target outreach to those workers to make them aware of their rights, Haffer said.

---

<sup>41</sup> Kehl, Danielle, Priscilla Guo, and Samuel Kessler. 2017. Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in sentencing. Responsive Communities Initiative, Berkman Klein Center for Internet & Society, Harvard Law School.

<sup>42</sup> Paige Smith, *Machine Learning Deployed to Help EEOC Predict Discrimination*, Bloomberg Law, Dec. 26, 2018, <https://news.bloomberglaw.com/daily-labor-report/machine-learning-deployed-to-help-eeoc-predict-discrimination>.