



October 28, 2016

Equal Employment Opportunity Commission  
Commission Meeting, EEOC Executive Officer  
131 M Street NE  
Washington, D.C. 20507  
commissionmeetingcomments@eoc.gov

Re: Big data and employment

To Chairwoman Yang:

The Center for Democracy & Technology (CDT) respectfully submits these comments in response to the Commission's October 13 public meeting on the implications of big data in employment. CDT is a nonprofit public interest organization dedicated to promoting digital privacy, free expression, and individual liberty. CDT commends the Commission's effort to protect equal opportunity in employment by examining how the collection and use of data may exacerbate structural inequality. In particular, we urge the Commission to scrutinize the use of automated decision-making systems in hiring, management, and employee evaluation practices and to ensure that such systems promote fairness, equality, and diversity.

### Introduction

Can algorithms deliver on their promises to hire more successful employees, improve office culture, and increase productivity? Responsible algorithms could be a force for good in the modern workplace, but the risks of harm that attend automated hiring, managing, and assessment systems are manifold. Rather than offering unbiased alternatives to human subjectivity, algorithms are imbued with the values of those who create them.<sup>1</sup> The natural inclination to hire people who look and act like us can be encoded and perpetuated by algorithms that seek to recreate historical patterns.<sup>2</sup> This bias toward past hires is likely to disproportionately favor historically advantaged groups—such as affluent, white, college educated males—to the detriment of historically disadvantaged groups.<sup>3</sup> Moreover, the trend toward collecting and analyzing large amounts of data on existing employees—whether to optimize performance or improve morale—threatens to erode important barriers between personal life and work life.<sup>4</sup> These comments

---

<sup>1</sup> See, e.g., A.R. Lange, Center for Democracy & Technology, *Digital Decisions: Policy Tools in Automated Decision-Making* 3–4 (Jan. 14, 2016), [https://cdt.org/files/2016/01/2016-01-14-Digital-Decisions\\_Policy-Tools-in-Auto2.pdf](https://cdt.org/files/2016/01/2016-01-14-Digital-Decisions_Policy-Tools-in-Auto2.pdf); Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Cal. L. Rev. 671, 677–80 (2016).

<sup>2</sup> Ali Lange, Katie McInnis & Michelle De Mooy, Center for Democracy & Technology, *Workplace Privacy: State Legislation & Future Technology Questions* 15 (June 1, 2016) [hereinafter *Workplace Privacy White Paper*].

<sup>3</sup> This also harms the economy and society at large, since considering a more diverse group of candidates increases the likelihood of finding the best person for the job and optimizing productivity.

<sup>4</sup> See, e.g., *Workplace Privacy White Paper* at 13–16.

1401 K Street NW, Suite 200 Washington, DC 20005

recommend steps the EEOC can take, and principles it can follow, in order to mitigate the negative consequences and support the positive impacts of big data and automated decision-making in the workplace.

### **Mitigating Negative Impacts**

The EEOC should provide guidelines for how algorithms must be designed, tested, implemented, and maintained by employers and any third parties they contract with in order to comply with existing employment and antidiscrimination laws. The Commission should also provide ethical guidelines that go beyond the letter of the law to promote diversity, inclusion, and fairness in employment. However, simply publishing and promoting best practices is insufficient. The EEOC should also perform audits of systems that use big data and algorithms to hire, manage, and evaluate employees. These audits should ensure first that the automated systems follow the law and additionally evaluate their compliance with the higher standards of diversity and inclusion promoted in EEOC best practices. Achieving fairness in automation will require that the EEOC consider the impact of automation on individuals systemically; not only evaluating each program in isolation, but considering how an individual experiences an automated hiring and employment experience from start to finish and how data flows between the different companies along the way.

#### Algorithms must be consistent with employment and antidiscrimination law

The Commission should audit hiring algorithms to ensure that they do not violate Title VII of the Civil Rights Act of 1964<sup>5</sup> or any other antidiscrimination law. As Solon Barocas and Andrew D. Selbst have explained, the goal of data mining algorithms is to rationally discriminate.<sup>6</sup> The objective of hiring algorithms is to determine a set of criteria by which to distinguish likely successful candidates from likely unsuccessful candidates. While few companies would expressly use race, sex, or age as a criterion, algorithms can still systematically discriminate based on these characteristics, often in unintentional and hidden ways. Certain criteria, such as credit, can act as proxies for sensitive characteristics.<sup>7</sup> With machine-learning algorithms—particularly those purchased from third parties—the employers themselves may not know which criteria the model is using to suggest new hires. Concrete metrics and guidance from the EEOC would make a tremendous difference by helping good actors live up to their values and creating a method for punishing those who do not.

#### Models should be based on sound research and logic

A predictive algorithm that can determine a person's livelihood should have a sound, statistically proven premise. Too often, predictive models are based on untested theories about how the world works, such as what makes a good employee.<sup>8</sup> Humans are naturally eager to explain, understand, and predict the chaos

---

<sup>5</sup> 42 U.S.C. 2000e et seq.

<sup>6</sup> Barocas & Selbst, *supra* note 1, at 677.

<sup>7</sup> Executive Office of the President, *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights* 13 (May 2016), [https://www.whitehouse.gov/sites/default/files/microsites/ostp/2016\\_0504\\_data\\_discrimination.pdf](https://www.whitehouse.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf).

<sup>8</sup> Laszlo Bock, *Here's Google's Secret to Hiring the Best People*, *Wired* (Apr. 7, 2015).

around them, but this hastiness to explain and predict complex phenomena risks oversimplification. For example, weather forecasting—while not perfect—is scientifically sound, because meteorologists have been systematically studying weather for many decades and fine tuning prediction models based on their objective scientific findings. These theories are falsifiable—they can be proven or disproven by other scientists.<sup>9</sup>

By contrast, the theories used to predict crime are unreliable. They are often based on the anecdotal observations of police officers and are not supported by strong statistical research. They are not falsifiable, because it is impossible to prove the error rate for negative predictions (that is, a prediction that says a crime *isn't* likely) and the results are skewed by feedback (police officers are more likely to observe crimes in neighborhoods where they increase patrols simply because they are spending more time there, not necessarily because they correctly predicted that a crime would occur there).

Hiring and managing employees resembles crime more than it does weather. Any number of factors, many of which are invisible, can contribute to making someone a good employee. There are no generally accepted, methodically tested, objective theories describing how to identify a good employee and, unlike with weather, the accuracy of predictions is not easy to track over time. Thus, the EEOC should discourage companies from relying too heavily on any particular set of criteria or trusting any model without human intervention.

Big-data models should proceed with caution when attempting to statistically predict subjective attributes such as personality, emotional state, or overall value as an employee

The EEOC should direct extra scrutiny toward predictive models that it cannot test. Highly subjective concepts, such as personality, emotion, and overall value as an employee, do not have broadly accepted definitions or proven sets of criteria. Because it is impossible to include every factor that might contribute to success, employers and programmers must make value judgments about which criteria to include. This type of design process makes it impossible to test the accuracy of a predictive algorithm, much less its inherent fairness. For example, an algorithm may accurately predict which employees have the criteria the algorithm is looking for—whether it be credit, longevity at a previous job, or proximity to the office—but that does not mean the algorithm is accurately predicting how successful employees will be.

Employers should not engage in any unnecessary data collection on, or tracking of, employees

Employers should collect only information that is necessary to effectively evaluate employee performance, and information collected for one purpose should not be used for a different purpose. Algorithms that evaluate workplace productivity and employee performance depend upon collecting information about employees. For example, Humanyze Workplace Solutions tracks employee movement and interaction through small chips in their ID cards. The premise of this tracking is to capture and analyze social interactions of

---

<sup>9</sup> See, e.g., Thomas M. Hamill et al., *Ensemble Reforecasting: Improving Medium-Range Forecast Skill Using Retrospective Forecasts*, 132 Monthly Weather Rev. 1434 (2004), [http://140.172.38.100/psd/people/tom.hamill/reforecast\\_mwr.pdf](http://140.172.38.100/psd/people/tom.hamill/reforecast_mwr.pdf).

employees with an eye toward understanding whether the company’s “top performers act differently.”<sup>10</sup> This type of tracking can chill interactions among employees and inhibit individual autonomy in the workplace.<sup>11</sup>

Employers should consider the context in which information was disclosed by the employee and/or collected by the employer before using it as input in an algorithm. Machine-learning tools can analyze employee communications across different internal platforms, such as workplace social networks and email, and attempt to discern and track employees’ “emotions” or “sentiments” from the text of their communications.<sup>12</sup> Even when such analyses are limited to workplace communications, aggregating and tracking communications and emotions across platforms and over time inhibits employees’ ability to maintain relative privacy. For example, if an employer uses sentiment analysis to analyze and keep track of an employee’s communications over time and notices that the employee’s communications often read as “sad,” the employer might infer—correctly or incorrectly—that the employee is depressed. Information about employee “sentiment” (which some might define as “morale”) might be considered—intentionally or unintentionally—in making decisions about whom to promote. Moreover, people use different platforms to communicate differently, so one algorithm may not be accurate with respect to multiple different platforms. For example, people tend to be more informal and include more personality in social media and chat conversations than in professional emails.

Research suggests that “continuous or unpredictable surveillance tends to lead to ‘more negative attitudes on the job and towards the organization’” and that “close monitoring leads employees to report more stress and feel they have no control organising their workloads.”<sup>13</sup> This calls into question the value of these technologies and whether the insights they reveal are worth the costs to overall job satisfaction.

Employee manuals should clearly state the employer’s policy with respect to the data it collects and algorithmically analyzes

The EEOC should require employers to disclose to employees the different types of information they collect about them and how they use this data. CDT has advocated for binding policies when it comes to the use and surveillance of employer-provided technology, such as cell phones.<sup>14</sup> The collection and use of big data is no

---

<sup>10</sup> Humanyze, <http://www.humanyze.com/index.html> (last accessed April 11, 2016) (“Visualize your team's engagement and cohesion communication network diagrams. Understand if your top performers act differently or if you have knowledge experts controlling communication flow in your organization.”).

<sup>11</sup> Workplace Privacy White Paper at 14.

<sup>12</sup> *Yammer Partners With Kanjoya to Offer Sentiment Analysis*, Market Wired (Aug. 2, 2012), <http://www.marketwired.com/press-release/yammer-partners-with-kanjoya-to-offer-sentiment-analysis-1686651.htm>; Kaveh Waddell, *The Algorithms That Tell Bosses How Employees Are Feeling*, The Atlantic (Sept. 29, 2016), <http://www.theatlantic.com/technology/archive/2016/09/the-algorithms-that-tell-bosses-how-employees-feel/502064/>).

<sup>13</sup> Nicole Kobie, *Big Brother Boss: The Psychological Weight of Workplace Monitoring*, ALPHR (Jan. 18, 2016), <http://www.alphr.com/business/1002466/bigbrotherbossthepsychologicalweightofworkplacemonitoring/page/0/1>.

<sup>14</sup> Alethea Lange & Katie McInnis, *With Workplace Privacy, Have a Policy and Follow the Policy*, Center for Democracy & Technology (Apr. 5, 2016), <https://cdt.org/blog/with-workplace-privacy-have-a-policy-and-follow-the-policy/>.

different. Along with following legal and ethical standards that might prohibit some data collection and use, employers should give clear and conspicuous notice to employees about data collection and analysis and should update policies whenever practices change.

### Supporting Positive Impacts

Data analytics have the potential to improve the workplace for everyone if used appropriately. These tools are best used to evaluate institutions, rather than people, to assess how managers and leadership can improve diversity, working conditions, and office culture. Instead of trying to predict how an individual will perform—an endeavor that is not well-supported by empirical research—algorithms are better suited to mining the past for insights into successful business practices and building sustainable relationships with employees.

#### Data analytics should be directed at the institution

Consider the example of sentiment analysis.<sup>15</sup> Twitter has used this tool to analyze responses to employee surveys.<sup>16</sup> The analysis looks for patterns in the responses, helping the company find common complaints and sentiments so that it can make more informed institutional changes.<sup>17</sup> This is a good example of using of data analytics to evaluate institutions in a way that does not violate expectations of privacy.

On the other hand, the same tool has been used to analyze employees' communications across workplace platforms and track their emotions over time, creating an emotional profile on each employee.<sup>18</sup> This type of tracking can lead to unintended privacy violations (e.g., discovering that someone suffers from depression) and inappropriate use of sensitive information to make decisions about hiring and promotion.

Algorithmic tools can be used for productive, privacy protective purposes that improve the workplace for everyone, or for invasive and unfair tracking of employees that crosses boundaries between work life and personal life. The EEOC should support the former and discourage the latter.

#### Good intentions aren't enough when using data to diversify hiring

Several startups have developed methodologies to help companies hire a more diverse workforce, but they don't all take the same approach with respect to their use of data. Comparing the approach of two

---

<sup>15</sup> See *supra* note 12.

<sup>16</sup> Waddell, *supra* note 12.

<sup>17</sup> *Id.*

<sup>18</sup> See *supra* note 12.

companies, Unitive<sup>19</sup> and Entelo,<sup>20</sup> demonstrates the variety of outcomes resulting from the same good intentions.

Unitive, a San Francisco based software company, works with companies to develop job postings that attract a range of candidates and subsequently helps employers “structure job interviews to focus on specific qualifications and mitigate the effect of the interviews' biases.”<sup>21</sup> Unitive’s methodology is grounded in social science research demonstrating that unconscious bias is real and nearly impossible to eliminate.<sup>22</sup> Their focus is on helping employers address and mitigate the effects of that bias in order to attract diverse candidates and give them a fair interview.

Entelo developed its products to solve the same problem, but takes a different approach. Entelo developed a platform that businesses can use to search for qualified women and minority job candidates,<sup>23</sup> much like a LinkedIn that includes diversity indicators. The platform uses data analytics to match candidates with employers seeking specific characteristics including gender, race or military considerations. The algorithm, “draws on information already publicly available, such as self-reported data on a social network like gender or affinity group memberships.”<sup>24</sup> This product gives employers working to increase the diversity of their workforce the opportunity to target their recruiting efforts at a guaranteed diverse pool.<sup>25</sup> But one side effect of this process is that potentially sensitive data is collected or inferred about individuals and tied specifically to their profile. This approach puts candidates in a position to prove or market themselves by sharing more information about who they are, and does so in the service of increasing diversity, rather than focusing on structural changes at the employer level.

These two tools are trying to solve the same problem, and they both rely on data analytics. But they operate under different philosophies—one is fundamentally about helping businesses alter their hiring behavior to reach diverse candidates and the other uses the blunt instrument of big data to fulfill the recruiting desires of companies. The second strategy doesn't necessarily address the larger cultural problems that affect retaining diverse candidates or reflect an understanding how homogeneity is created in the first place. The EEOC must be aware that good intent is not sufficient in determining the quality of the product.

## Conclusion

---

<sup>19</sup> <http://www.unitive.works/>.

<sup>20</sup> <https://www.entelo.com/>.

<sup>21</sup> Yuki Noguchi, *How Startups Are Using Tech to Try and Fight Workplace Bias*, NPR (Sept. 1, 2015), <http://www.wnyc.org/story/how-startups-are-using-tech-to-mitigate-workplace-bias/>.

<sup>22</sup> *Id.* See also, e.g., Project Implicit, <https://implicit.harvard.edu/implicit/takeatest.html>.

<sup>23</sup> Entelo, <https://www.entelo.com/>; Cat Zackrzewski, *Entelo Wants to Help Companies Hire Diverse Employees as They Scale Up*, Tech Crunch (March 19, 2015), <https://techcrunch.com/2015/03/19/entelo-wants-to-help-companies-hire-diverse-employees-as-they-scale-up/>.

<sup>24</sup> Zackrzewski, *supra* note 23.

<sup>25</sup> Entelo, *Entelo Diversity*, <https://www.entelo.com/outbound/diversity/>

Bias in data-driven systems is nearly unavoidable.<sup>26</sup> But that doesn't mean that these tools should not play a role in our workplace, and in fact it seems certain that tools that assist human decision makers are here to stay. Rather than working toward demanding unbiased programming or data, regulators and institutions should focus on regulating the role these programs can play in institutional decisions and consider incentives that encourage scrutiny turned to employers rather than employees.

The EEOC is well positioned to develop best practices, audit mechanisms, and basic rules that ensure that civil rights law is not violated and that democratic values are embedded into these technologies. Audits will reveal weaknesses in programs and assumptions that can be mitigated technically and/or factored into the decisions made by individuals using the outputs of the data. This will require the agency to take a holistic approach to hiring and employment automation and analytics, not only evaluating programs or companies in isolation but considering how they interact with each other as data is shared or sold, and assumptions made by one hiring program might be carried forward into another. For example, an individual who is tagged by an automated hiring system as undesirable in some way, or just not a good fit, with no transparency or opportunity for redress, may be permanently damaged if this industry is allowed to grow, unchecked by government regulators.

Thank you again for the opportunity to submit comments, and we welcome any questions or comments.

Sincerely,

Natasha Duarte  
Ron Plesser Fellow  
Center for Democracy & Technology

Alethea Lange  
Senior Policy Analyst  
Center for Democracy & Technology

---

<sup>26</sup> Moritz Hardt, *How Big Data is Unfair*, Medium (Sept. 26, 2014), <https://medium.com/@mrtz/how-big-data-is-unfair-9aa544d739de#.aug724ksm>.