RISE OF THE MACHINES
THE BENEFITS AND RISKS OF THE USE OF AI IN EMPLOYMENT DECISIONS

BY CHRISTINE M. SNYDER & ARIANA E. BERNARD

In a world where parsing through employment applications can feel like "finding a needle in a haystack," employers are rapidly turning to artificial intelligence (AI) for assistance. While this does not signal the dawn of an age where robots completely take over the workforce, it does mean that AI has gained a foothold in the employment decision-making process.

The use of AI has many benefits, including increased efficiency in screening applicants, but reliance on AI also presents concerns that the algorithms used to do so could run afoul of federal anti-discrimination protections like Title VII of the Civil Rights Act and the Americans with Disabilities Act and similar state and local laws.

The rapid expansion of AI in employment decisions has fostered vigorous debate regarding the relative benefits and risks of AI, which has largely gone unregulated thus far. That may soon change. The Equal Employment Opportunity Commission (EEOC) has launched an agency-wide initiative to ensure that the use of AI complies with federal civil rights laws, and as part of that initiative, it held a public hearing on January 31, 2023, drawing testimony from a variety of experts in law, civil rights, computer science, data analytics, and employee and employer advocacy.

One thing is clear — federal, state, and local governments are beginning to take notice of the use of AI in employment decisions, and the first moves to regulate its use are underway.

What Is AI and How Is It Used by Employers?
Generally, in the employment context, AI is the use of machine-based systems to make predictions, recommendations, or decisions based on a set of criteria to achieve a given objective. Some AI statically applies fixed, objective criteria, like education level, while other AI systems apply more abstract concepts into models in an automated manner (i.e., machine learning). The most familiar use of AI in employment decisions is to sort and evaluate job applicants to identify those best suited for interview or hire.

Employers are increasingly turning to AI. According to Chair Charlotte Burrows of the EEOC, upwards of 83% of employers use some form of automated tool to screen or rank candidates. Employers utilize AI to assist in a wide range of employment practices. At the recruiting and hiring stage, AI can be used to target job ads, analyze resumes, and even interview candidates, where AI can be used to analyze facial expressions, eye contact,

> The rapid expansion of AI in employment decisions has fostered vigorous debate regarding the relative benefits and risks of AI, which has largely gone unregulated thus far. That may soon change.

and word choice to assess them. AI can also be used to monitor employee productivity and performance, which can then serve as the basis for other employment decisions like promotion, salary, and termination.

The Case for AI: AI Can Mitigate Bias and Increase Diversity
Proponents of AI tout how it benefits employers by significantly reducing costs and time spent on recruiting and hiring. More importantly, proponents argue that AI can foster a more diverse workforce and reduce inherent human biases through the use of selection algorithms that are blind to a candidate's age, race, gender, disability, or other protected characteristic.

Such biases are often unconscious, but their effect can be significant, as demonstrated in a 2005 Yale University study of the evaluation of candidates for police chief. In that study, human evaluators routinely justified choosing men over women, shifting their preferred criteria in order to do so. Men without college degrees were chosen over women with college degrees on the basis that "street smarts" were more important. However, when applicant names were swapped, evaluators chose men with college degrees over women without college degrees, shifting their justification to value education more highly. (*See* Eric Luis Uhlmann & Geoffrey L. Cohen, *Constructed Criteria Redefining Merit to Justify Discrimination*, 16 Psychol. Sci. 474, 474 (2005)). Proponents of AI argue that its use can significantly reduce such unconscious bias.

The Case for Caution: AI Can Perpetuate and Amplify Discrimination
Despite its potential benefits, AI also poses risks of greater discriminatory outcomes. Algorithms that rely on existing employee data may perpetuate existing biases. For example, if a company's baseline data is its existing workforce, which already lacks diversity, then the AI's profile of a potentially successful applicant may embed preexisting biases into the AI selection criteria.

Further, many seemingly innocuous data points can serve as proxies for those characteristics. For example, while an employer may simply want to identify candidates close to the office, screening for zip codes can often unintentionally result in racial bias. Likewise, screening out applicants with gaps in work history can disproportionately screen female applicants who more regularly take time off to care for children or family members. Further, some evidence indicates that AI that analyzes facial expressions and word choice disproportionately screens out minority candidates.

Additionally, even if input data is carefully crafted, AI creates risks of feedback loops that may lead to bias. For example, if older candidates interact poorly or more slowly with an AI system, the algorithm may "learn" to disfavor similar candidates.

Essentially, algorithms that rely on incomplete, unrepresentative, or error-ridden data are likely to lead to biased outcomes and expose employers to claims of discrimination under federal, state, and local anti-discrimination laws. Admittedly, it is unlikely that employers will instruct AI algorithms to intentionally reject individuals in protected classes. More likely, employers could face disparate impact claims when they rely on facially neutral factors, but the results tend to reject members of protected classes at higher rates.

The Need for Guidance and a Balanced Approach

Currently, there are no federal regulations governing employers' use of AI in employment decisions. There have been steps, however, to delve into the issue. For example, the White House Office of Science and Technology Policy has created a Blueprint for an AI Bill of Rights, which identifies principles to guide the design, use, and deployment of AI. Also, several states have enacted statutes regulating the use of AI. For example, Maryland and Illinois have added applicant notice and consent requirements. Similarly, California has proposed regulations to clarify that its anti-discrimination laws also apply to the use of AI, and New York City's AI regulations require independent bias audits of AI and the publication of audit summaries.

As the use of AI in the employment space rapidly expands, employers should expect that more regulations and guidance are to come. It remains unanswered what form they might take, but the testimony from a variety of panelists during the January EEOC hearing identified several ways that the EEOC could institute requirements to mitigate potential bias in the use of AI. The following are the most common suggestions to come out of the hearing and that are likely to be reflected in anticipated regulations and guidance:

- Transparency – Any effort to regulate the use of AI in workplace decisions will likely include a requirement that employers disclose their use of AI to applicants and employees.
- Vendor Vetting – Since most employers rely on AI products and services from third-party vendors, employers will likely need to demonstrate it properly vetted vendors and their products. To aid in this effort, the non-profit Data & Trust Alliance has developed a set of "safeguards" for the use of AI, including a list of 55 questions for employers to ask AI vendors.
- Auditing and Reporting – Many experts who testified before the EEOC discussed requiring or recommending that employers conduct audits to detect and remove potential bias from their AI systems. It remains to be seen what standards would apply to such audits, but employers should anticipate that they will need to periodically audit their AI-generated results, document or report those findings, and demonstrate compliance with yet-to-be-defined parameters.

One thing is for certain: AI is here to stay, and while it has many benefits, it also brings significant risks that, if not implemented properly, can perpetuate and intensify workplace discrimination. Right now, a growing number of employers, employees, and candidates are using or are subject to AI, but in the absence of any meaningful regulation or guidance on its proper use. As regulators look to establish safeguards, they should carefully balance the need to protect against unlawful discrimination, but also not completely discourage the use of AI and stifle its potential benefits. Regardless, the largely unchecked use of AI in the employment context is likely coming to an end.

Christine Snyder is a partner at Tucker Ellis LLP in the firm's Labor and Employment Group where she assists employers navigate complex employment issues. She has been a member of the CMBA since 2013. She can be reached at (216) 696-5593 or christine.snyder@ tuckerellis.com.

Ariana Bernard is an associate at Tucker Ellis LLP. Her practice includes a focus on employment litigation where she defends employers in a wide variety of employment cases. She has been a member of the CMBA since 2020. She can be reached at (216) 696-2965 or ariana.bernard@ tuckerellis.com.