

Confronting the Threat of Employment Discrimination in the Algorithmic Age

Before the Equal Employment Opportunity Commission

January 31, 2023

**Testimony of Matthew U. Scherer, Senior Policy Counsel, Workers' Rights and
Technology
Center for Democracy & Technology**

Chair Burrows, Vice Chair Samuels, and Commissioners,

Thank you for the opportunity to testify on employment discrimination in A.I. and automated systems. My name is Matt Scherer, and I am Senior Policy Counsel for Workers' Rights and Technology at the Center for Democracy & Technology. CDT is a nonprofit, nonpartisan 501(c)(3) organization based in Washington, D.C. that advocates for stronger civil rights protections in the digital age. CDT's work includes a project focused on how algorithmic tools that are used in employment decisions can interfere with workers' access to employment and limit their advancement opportunities.¹

CDT has worked with a broad coalition of civil rights and civil society organizations over the past several years to help develop principles and standards regarding these technologies that center and advance the interests of workers, particularly those from historically marginalized and disadvantaged groups. In particular, over the past two years, we worked with several of these organizations to create the *Civil Rights Standards for 21st Century Employment Selection Procedures*.² We were proud to partner with the ACLU, American Association of People with Disabilities, Upturn, the Leadership Conference for Civil and Human Rights, and the National Women's Law Center in drafting the principles, and to receive endorsements from the Bazelon Center for Mental Health Law, Color of Change, the National Employment Law Project, the Autistic Self-Advocacy Network, and other groups.

As the Commission is aware, more and more employers are using artificial intelligence and other automated systems to make employment decisions that determine the course of workers'

¹ <https://cdt.org/area-of-focus/privacy-data/workers-rights/>.

² Ctr. for Democracy & Tech., et al., *Civil Rights Standards for 21st Century Employment Selection Procedures*, <https://cdt.org/insights/civil-rights-standards-for-21st-century-employment-selection-procedures/>.

careers and lives. Automated employment decision tools (AEDTs) come in many forms, including tools that analyze the words candidates use in resumes and programs that use computer games or quizzes to estimate a candidate's personality traits.

But these tools rarely, if ever, make an effort to directly measure a worker's ability to perform the essential duties and tasks that will be expected of whomever the employer hires for the position. They also often pose a risk of discrimination against already-disadvantaged groups of workers, who are often underrepresented in the data used to train employment decision tools and whose relevant skills and abilities may not be as obvious to an automated system. My testimony will discuss how the current legal framework fails to adequately account for the unique risks of discrimination that AEDTs present and discuss how the Civil Rights Standards are a key resource that the Commission should use to inform future guidance, technical assistance, and regulatory efforts.

The Current Legal Framework Does Not Adequately Address the Heightened Discrimination Risks That AEDTs Pose

From a civil rights perspective, the current legal landscape governing AEDTs needs clarification and refinement. While the Uniform Guidelines for Employee Selection Procedures (UGESPs)³ remain in effect, they do not adequately reflect the many changes in law and social science that have occurred in the five decades since they were drafted.

The Commission and its sister agencies adopted the UGESPs in 1978, more than a decade before Congress passed the Americans with Disabilities Act (1990). By their own terms, the UGESPs do not address discrimination against people with disabilities or age discrimination, nor do they address the full scope of sex discrimination. The UGESPs also have not been updated to expressly incorporate modern scientific standards regarding validation and fairness.⁴ This makes further action by the EEOC urgent, to clarify how the EEOC will interpret and apply the statutes and regulations it enforces to meet the unique risks posed by automated tools.

To begin, Title VII (and the ADA) state that where an employment practice has a disparate impact, it constitutes unlawful discrimination unless the employer demonstrates "that the practice is job related *for the position in question* and consistent with business necessity."⁵ The phrase "for the position in question" means that for a test that has a disparate impact to be valid,

³ 29 C.F.R. § 1607.1 et seq.

⁴ See generally Am. Educ. Research Ass'n, et al., *Standards for Educational and Psychological Testing*, (4th ed. 2014); Soc'y for Indus. & Organizational Psychology, *Principles for the Validation and Use of Personnel Selection Procedures* (5th ed. 2018). See also Matthew U. Scherer, et al., *Applying Old Rules to New Tools: Employment Discrimination Law in the Age of Algorithms*, 71 So. Car. L. Rev. 449 (2019), available at https://ssrn.com/abstract_id=3472805.

⁵ 42 U.S.C. § 2000e-2(k)(1)(A)(i) (Title VII) and 42 U.S.C. § 12112(b)(6) (ADA) (emphasis added).

an employer must link it to the duties *of the specific job for which it is being used*. This echoes the Supreme Court’s admonition in *Griggs v. Duke Power Co.* that any test or screening mechanism for job applicants “must measure the person for the job and not the person in the abstract” to survive a Title VII challenge.⁶

Many of the AEDTs being marketed today fail to meet this job-specific validity requirement because they have one (or both) of two characteristics: (1) they measure abstract or amorphous characteristics not tailored to the job in question; and/or (2) they rely on machine-learning techniques that use correlation alone—rather than a logical or causal relationship with job functions—to establish a link between test results and job performance.

Tools That Measure Abstract Candidate Characteristics

Tools offered by some of the more prominent AEDT vendors claim to rate candidates not on specific knowledge or abilities, but on highly abstract and subjective qualities like “empathy,” “influence,” and “personality.”⁷ CDT’s 2020 report, *Algorithm-driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?*, describes in detail how these tools can discriminate against workers based on attributes including race, sex, national origin, and disability status. When such tests result in disparate impacts or tend to screen out disabled workers, federal law requires employers to establish job-relatedness in order to survive a discrimination claim. That showing is not tenable with tools that purport to measure generic personality traits or other characteristics untethered from the specific duties or essential functions of the jobs for which candidates are being evaluated. That is precisely the sort of measurement of “the person in the abstract,” rather than for a specific job, that the Supreme Court warned against and that the text of Title VII and the ADA expressly prohibits.⁸

The guidance document that the EEOC published last year on AEDTs and the ADA is a great first step in pushing back against the use of such tests. It recognizes that to minimize the risk of unlawful disability discrimination, employers should ensure that “tools only measure abilities or qualifications that are truly necessary for the job—even for people who are entitled to an on-the-job reasonable accommodation.”⁹ I encourage the full Commission to take the next step

⁶ *Griggs v. Duke Power Co.*, 401 U.S. 424, 436 (1971).

⁷ See, e.g., Matthew Scherer, HireVue “AI Explainability Statement” Mostly Fails to Explain What it Does (2022), https://cdt.org/insights/hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does/?utm_source=rss&utm_medium=rss&utm_campaign=hirevue-ai-explainability-statement-mostly-fails-to-explain-what-it-does (noting how the competencies that one vendors’ assessments claim to measure “are not moored to the actual responsibilities and functions of specific jobs”).

⁸ 42 U.S.C. § 2000e-2(k)(1)(A)(i); 42 U.S.C. § 12112(b)(6).

⁹ EEOC, *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees*, Q14 (Promising Practices for Employers), <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence#Q14>.

by elevating this from a best-practice recommendation to a rule that the EEOC will enforce. I respectfully submit that the plain terms of our antidiscrimination laws require nothing less.

Tools That Rely Exclusively on Correlation

Another problem with many AEDTs stems from the fact that machine-learning algorithms do not examine whether the attributes that a model uses to predict job performance are logically or causally related to the essential functions of a job, nor do they analyze whether the attributes in the training data include a set of variables that are representative of the skills needed to perform a particular job.¹⁰ Instead, AEDTs that rely on such algorithms depend on correlation—which, in the employment selection processes, means the degree to which differences in “predictor” attributes that could be used to assess candidates (such as years of experience or schools attended) are associated with differences in some target “criterion” connected to job performance (such as supervisory ratings or sales figures). Using correlation alone to select which “predictor” variables an AEDT will use can lead to both invalidity and discrimination for two reasons.

First, a tool built on correlation-based techniques alone is highly unlikely to capture all (or a representative set) of the essential functions of a specific job. Say a company wants to screen marketing specialists using a resume scanning algorithm that uses machine learning to decide which words in resumes are predictive of job performance. Even if thousands of qualified marketing specialists’ resumes are included in the algorithm’s training data, there are aspects to marketing jobs (such as interpersonal skills) that cannot be reliably extracted from a resume alone—and the fact that certain terms often show up in the resumes of successful workers does not necessarily mean those terms are the best indicators of a person’s ability to do the functions of a job (as discussed further below).

Few, if any, data sets are rich enough to cover all of the essential knowledge and abilities needed for a given job, much less the nuances of how such knowledge and abilities will be needed for a role at a specific company. This means that any tool that operates solely by searching for correlations in historical data sets will create an incomplete picture of a candidate’s ability to perform the job in question.¹¹ If a company then over-relies on such a tool when making employment decisions, its decisions will not be based on an adequate representation of essential job functions, as both the ADA and the UGESP require.

Second, the use of correlation-driven statistical methods increases the risk that a tool will capitalize on correlations that are due to chance rather than due to a logical, causal, or organic

¹⁰ See Scherer et al., *supra* note 4, at 487.

¹¹ This shortcoming is referred to as *construct deficiency* in the literature on test validity. That is, the test does not capture all of the relevant aspects of the construct (in this case, the ability to perform essential job functions) that the tool is supposed to be measuring.

relationship with job performance. As a result, AEDTs may discover and use attributes that are *construct-irrelevant*—that is, attributes that are tied not to job-performance factors (the “construct” that employment tests are supposed to capture), but to irrelevant characteristics.¹² This can lead to differences in scores or selection decisions that are due to factors unrelated to a candidate’s ability to perform essential job functions.¹³ This can happen, for example, when a test measures something more or different than the relevant aspects of job performance (e.g., if a test of oral communication skills is affected by a test-taker’s proficiency in written English); or when outcomes reflect cultural differences rather than (or in addition to) differences in job-related competencies.

A hypothetical example from an article I co-authored illustrates how this can lead to unlawful discrimination:

[S]ay that a company was training an algorithmic tool to recognize good software engineers using training data that reflects the demographics of their best current network engineers, who are predominantly white males. If these employees share, as is likely, construct-irrelevant characteristics that are reflected in the training data, the tool will learn to associate those characteristics with good job performance. This could have two related adverse impacts on qualified candidates who are not white males. First, if the ablest female and nonwhite candidates have attributes (whether construct-relevant or not) that differ from those of the white males who dominate the current sample, the tool’s accuracy will be lower when scoring those candidates, just as the gender classification programs in the MIT [Gender Shades] study^[14] were less accurate when attempting to classify individuals with darker skin. Second, the individuals that the tool identifies as the best candidates from the underrepresented groups may have scored highly not because of characteristics that affect their actual competence, but because of the construct-irrelevant characteristics they share with the current software engineers.

Both of those factors may drive down the number of qualified female and minority candidates that the tool selects. In addition, the candidates who the tool does recommend from the disadvantaged group are less likely to be the most

¹² See Keith E. Sonderling, et al., *The Promise and The Peril: Artificial Intelligence and Employment Discrimination*, 77 U. Miami L. Rev. 1, 24 (“In analyzing a large quantity of data, an algorithm might identify a statistical correlation between a specific characteristic of a job applicant and future job success that nevertheless lacks a causal relationship.”).

¹³ The technical term for this phenomenon is *construct-irrelevant variance* or *contamination*.

¹⁴ In that study, facial recognition technology was found to be less accurate in correctly identifying the gender of darker-skinned individuals than lighter-skinned ones—and the darker the individual’s skin, the less accurate the tool was.

competent candidates from that group, which may reduce the likelihood that they are ultimately hired and retained. Through these mechanisms, an employer's adoption of an algorithmic tool could inadvertently reinforce existing demographics.¹⁵

When adverse impacts arise because members of a group perform differently on improperly included or excluded aspects of job performance, the resulting discriminatory impacts would constitute Title VII and ADA violations.¹⁶

The UGESPs, having been written long before the rise of machine-learning algorithms that can comb through hundreds or thousands of potential predictors and build a model based solely on correlation, do not adequately address this source of discrimination and invalidity in AEDTs. Again, the ADA guidance that the EEOC published last year is encouraging in this regard. That guidance suggests that employers ensure that “necessary abilities or qualifications are measured directly, rather than by way of characteristics or scores that are correlated with those abilities or qualifications.”¹⁷

Here too, I encourage the full Commission to take further formal action to stop the proliferation of discriminatory tools that rely on aimless correlations. The Commission should issue additional guidance explaining that, consistent with the proper interpretation of “job-relatedness” under our antidiscrimination laws and *Griggs*, correlation alone does not suffice to establish job-relatedness absent other evidence or explanation addressing *why* the attributes measured by an automated tool should be expected to predict a candidate's ability to perform essential functions. Additionally, the EEOC should issue guidance and, if practicable, engage in rulemaking to address how impact and validity analyses should be conducted in light of the unique requirements of the ADA—a statute to which the UGESPs do not apply and that thus presents an important area for agency action.

Civil Rights Standards for 21st Century Employment Selection Procedures

Despite the threats to validity and the risk of discrimination that AEDTs pose, some vendors and allied special interest groups have actively sought policy changes that would weaken or undercut existing protections or confuse employers and consumers about what current law requires. They often do so under the pretense that their technologies are less biased than longer-established employment decision processes, and that their proposed policy changes

¹⁵ Scherer, *supra* note 4, at 488.

¹⁶ See 42 U.S.C. § 2000e-2(k)(1)(A)(i) (demonstration of job-relatedness required to overcome showing of adverse impact).

¹⁷ EEOC ADA Guidance, *supra* note 11.

thus represent a pro-civil-rights position. The evidence and arguments used to support these efforts are generally incomplete at best, and highly misleading at worst.¹⁸

Faced with this combination of (i) the risks of wide-scale discrimination posed by AEDTs and (ii) intensifying efforts to insulate AEDTs from discrimination accountability, CDT partnered with many of the nation’s leading civil rights organizations—including the ACLU, which is here today—to create the *Civil Rights Standards*.¹⁹ While the rise of automated tools was the impetus for the *Standards*, the *Standards* themselves apply to all formalized selection procedures and thus lay out a path to updating existing rules and guidance regarding employee assessments.

It is our hope that the Commission considers the *Standards* as it completes its Strategic Enforcement Plan, and more generally as it moves forward in its regulatory, educational, and compliance efforts regarding automated tools.

The *Standards* advance the five *Civil Rights Principles for Hiring Assessment Technologies*²⁰ that were first developed in 2020 by a broad coalition of civil rights groups, including CDT. Those five principles are: Nondiscrimination, job-relatedness, auditing, notice and explanation, and oversight and accountability. The *Standards* expand on and operationalize these five Principles, providing a concrete alternative to proposals that would set very weak notice, audit, and fairness standards for automated tools. They are designed for inclusion in regulatory guidance, for adoption by vendors and companies, and for workers who deserve to know their rights.

The *Standards*’ key provisions include:

Nondiscrimination

Targeting and reducing the risk of all forms of unlawful discrimination by:

- Requiring companies to take a proactive approach to eliminating potential sources of discrimination

¹⁸ See, e.g., Scherer, *supra* note 7; Hilke Schellmann, *Auditors are testing hiring algorithms for bias, but there’s no easy fix*, MIT Technology Review, Feb. 11, 2021, <https://www.technologyreview.com/2021/02/11/1017955/auditors-testing-ai-hiring-algorithms-bias-big-questions-remain/>; Mona Sloane, *The Algorithmic Auditing Trap*, OneZero (Medium), Mar. 16, 2021, <https://onezero.medium.com/the-algorithmic-auditing-trap-9a6f2d4d461d>.

¹⁹ As of January 13, 2023, the following organizations have endorsed the *Standards*: Center for Democracy & Technology (CDT), American Association for People with Disabilities (AAPD), American Civil Liberties Union (ACLU), Autistic People of Color Fund, Autistic Self Advocacy Network (ASAN), Autistic Women & Nonbinary Network (AWN), Bazelon Center for Mental Health Law, Color Of Change, The Leadership Conference on Civil and Human Rights, National Employment Law Project (NELP), National Women’s Law Center (NWLC), TechEquity Collaborative, Upturn.

²⁰ *Civil Rights Principles for Hiring Assessment Technologies* (2020), <https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/>.

- Mandating that employers use the least discriminatory selection procedure (SP) available
- Banning certain SPs that pose an especially high risk of discrimination

Job-Relatedness

Ensuring that SPs only measure traits and skills that are important to job performance by:

- Requiring SPs to measure only the essential functions of the job(s) for which they are used
- Requiring audits to include a description of the essential functions for which the SP is being used and an explanation of why those functions are essential to the position
- Requiring correlation-based evidence of validity to be supported by theoretical, logical, or causal reasoning sufficient to explain *why* the SP's predictors should be expected to predict the ability to perform essential functions
- Prohibiting the use of SPs whose validity cannot be assessed

Auditing

Requiring both *pre-deployment* and *ongoing audits* by an independent auditor. Audits must:

- Include a thorough job analysis for each position for which the SP would be used
- Analyze the SP's validity and risk of various forms of discrimination
- Determine whether valid and less discriminatory assessment methods are available

Notice and Explanation

Creating three layers of disclosure requirements, each tailored to a different intended audience:

- A short-form disclosure describing for candidates how the SP works and how they can raise concerns and access accommodations
- Detailed audit summaries, intended for sophisticated stakeholders like regulators and workers' advocates
- Comprehensive recordkeeping obligations so all information is preserved in case of investigation or litigation

Oversight and Accountability

Giving all stakeholders a role in ensuring that selection procedures do not violate the law, by:

- Making it unlawful to use or sell discriminatory SPs
- Giving candidates the right to communicate concerns about SPs prior to their use, and the right to request human review of automated SPs' recommendations
- Providing for agency enforcement, as well as a private right of action for certain unlawful practices
- Making employers and vendors jointly responsible for audits, and jointly and severally liable for discrimination

The *Civil Rights Standards* provide a roadmap to managing the risks associated with modern selection tools while centering the rights and dignity of workers, particularly those most vulnerable to marginalization and discrimination. They contain provisions that would address the

unique threats of discrimination discussed above. They are designed to be modular; each standard reinforces and strengthens the others, but each also stands on its own. Again, we hope they can be a resource to the Commission as it completes its Strategic Enforcement Plan and continues its important regulatory, educational, and compliance efforts regarding automated tools.

Conclusion

We have seen many ways in which new technologies have made the workplace and labor market fairer and more efficient. The rise of the Internet, for example, enhanced workers' ability to search and apply for jobs and career paths. But not all new technologies represent progress. History is replete with examples of supposed innovations that, despite the hype and assurances of the companies promoting them, failed to live up to their potential or were rushed to market before the technology was ready for prime time. Where, as here, the careers and livelihoods of so many workers are at stake, there is a risk of great harm if ineffective and potentially harmful technologies are allowed to proliferate without proper scrutiny.

As the Commission completes its Strategic Enforcement Plan over the coming weeks and months, we urge it to use its platform and authority to ensure that workers, not machines, remain at the center of the future labor market. The rights of workers, particularly vulnerable and marginalized workers, must not be trampled or glossed over for the sake of convenience or efficiency. Thank you.