

The Guidance Is Gone but the Statute Is Not: AI, Disability, and Employment Decisions After the January 2025 Withdrawals

Heather Grizzle · April 2, 2025

The EEOC withdrew its AI guidance in January 2025. The ADA did not change. What employers still owe disabled applicants under automated hiring assessment.

Contents

1. What changed, and what did not
2. Why disability exposure does not behave like bias exposure
3. What the March 2025 complaint shows
4. Accessibility, accommodation, and validity are three different obligations
5. The state layer and its timetable
6. What to do now
7. Conclusion
8. References

On January 23, 2025, Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence, revoked the prior administration's AI executive order and directed federal agencies to review and roll back AI-related policy. Within days, the practical consequences reached the employment context.

What changed, and what did not

On January 27, the Equal Employment Opportunity Commission removed its artificial intelligence guidance from its website, including the May 2022 technical assistance document on the Americans with Disabilities Act and the use of software, algorithms, and artificial intelligence to assess job applicants and employees, its companion Tips for Workers summary, and the May 2023 technical assistance document addressing Title VII. The Office of Federal Contract Compliance Programs removed its guidance on artificial intelligence and equal employment opportunity for federal contractors. The Department of Labor noted that the AI and Inclusive Hiring Framework published in September 2024 by the Partnership on Employment and Accessible Technology, and its own October 2024 artificial intelligence best practices, may no longer reflect current policy.

A great deal of commentary has read these removals as a reduction in employer exposure. That reading confuses guidance with law. Technical assistance documents are interpretive. They explain how an agency understands a statute and how it intends to exercise its enforcement discretion. They do not create the obligation and their withdrawal does not extinguish it. The Americans with Disabilities Act was enacted in 1990 and has not been amended. Title I still prohibits discrimination on the basis of disability in the terms and conditions of employment, still requires reasonable accommodation absent undue hardship, still prohibits the use of qualification standards and selection criteria that screen out individuals with disabilities unless the criterion is job-related and consistent with business necessity, and still restricts disability-related inquiries and medical examinations before a conditional offer. The Uniform Guidelines on Employee Selection Procedures at 29 C.F.R. Part 1607 still direct users of selection procedures to maintain records of impact and, where impact exists, evidence of validity. The recordkeeping regulations at 29 C.F.R. Part 1602 still require preservation of application records and, once a charge is filed, preservation of relevant records until final disposition.

What the January withdrawals removed was not liability. It was notice. Employers lost the document that told them, in plain language, which of their practices the agency considered problematic and why. The exposure remains and the map is gone.

Why disability exposure does not behave like bias exposure

The exposure remains and the map is gone.

The dominant response to algorithmic risk in hiring is the bias audit, and the dominant metric is the four-fifths rule, which compares selection rates across groups and treats a ratio below eighty percent as evidence of adverse impact. New York City's Local Law 144, in effect since July 2023, requires an annual independent bias audit of automated employment decision tools and public posting of the results. Vendors increasingly market audit results as assurance.

This apparatus is poorly suited to disability, and the mismatch is structural rather than a matter of audit quality. Group-statistical impact testing requires stable, known, mutually exclusive categories with enough members in each to support a rate comparison. Race and sex approximate those conditions. Disability does not. It is heterogeneous to the point that a single category has little analytic meaning, since the accommodation needs and interaction patterns of a Deaf applicant, a blind applicant, an applicant with a motor impairment, and an applicant with a psychiatric disability have almost nothing in common. It is frequently undisclosed, because applicants have no obligation to identify themselves and strong incentives not to. It is often low-prevalence in any single requisition, which leaves subgroup comparisons underpowered even where identification exists. And under the ADA the operative unit of analysis is the individual rather than the group, because a screen-out claim asks whether a criterion tends to exclude an individual with a disability, and a failure-to-accommodate claim asks what happened when one person made one

request.

The consequence deserves stating plainly. An automated selection system can pass every bias audit conducted on it, publish an impact ratio comfortably above the four-fifths threshold on every category measured, and still be excluding disabled applicants one at a time in a manner that is individually actionable and statistically invisible. Audit results are not evidence of ADA compliance. Presenting them as such, which is now common vendor practice, creates a documented record that the employer relied on an assurance that did not address the risk.

This is also why the widely repeated claim that well-designed algorithms reduce human bias, and therefore level the playing field for disabled candidates, should be treated as unevidenced. It may be true for some forms of bias. There is no public evidence base establishing it for disability, and the mechanism by which these systems operate suggests the opposite. Models score conformity to a distribution learned from prior successful candidates. Disability, in the dimensions these systems most often measure, is precisely what nonconformity to that distribution looks like.

What the March 2025 complaint shows

On March 19, 2025, the American Civil Liberties Union, the ACLU of Colorado, Public Justice, and Eisenberg and Baum LLP filed a complaint with the Colorado Civil Rights Division and the EEOC on behalf of a Deaf and Indigenous woman who was required to complete an automated video interview when she applied for a promotion. The complaint alleges violations of the ADA, Title VII, and the Colorado Anti-Discrimination Act. According to the public filings, she requested human-generated captioning as a reasonable accommodation, the request was denied on the basis that the platform's automated captions would suffice, portions of the interview were nonetheless without captions, and she was subsequently rejected on grounds concerning communication. The employer states that the allegations

are without merit and that it provides reasonable accommodations to all candidates. The vendor states that the complaint is without merit and rests on an inaccurate assumption about the technology used. These are allegations and denials, and an agency investigation will determine whether cause exists.

Regardless of outcome, the fact pattern is the one employers should be modeling, because it does not depend on proving that a model is biased. It depends on three ordinary questions. Was the assessment accessible to this candidate. Was an accommodation requested, and what was the response. Can the employer explain what the assessment measured and why that construct is job-related.

Note also what the theory does not require. It requires no comparative statistics, no expert reconstruction of a model, and no disparate impact analysis. The evidence is a dated accommodation request, an answer to it, and an adverse outcome. That evidentiary profile is inexpensive to assemble and it is created by the employer's own records.

Accessibility, accommodation, and validity are three different obligations

Much of the current discussion collapses distinctions that operate independently, and conflating them produces compliance programs that address one obligation while leaving two open.

Accessibility is a property of the interface. Whether the assessment platform can be operated with a screen reader, whether media is captioned, and whether timing can be adjusted are questions about whether the candidate can reach the task at all. This is the layer most organizations have addressed, usually through a vendor conformance claim.

Accommodation is a process obligation. It requires that candidates be told before the assessment begins that modifications are available, that a request triggers an interactive

process rather than a form rejection, that a decision-maker is identified, and that an alternative pathway exists which does not itself disadvantage the person who used it. A platform that is technically accessible in every respect can still generate a failure-to-accommodate claim if no one told the applicant that an alternative existed.

Validity is a property of the measurement. It asks whether what the system scores corresponds to a requirement of the job. This is the layer almost nobody has addressed, and it is where automated video and speech assessment is most exposed. Where a system scores fluency, articulation rate, response latency, prosody, facial affect, or a construct labeled communication or engagement, the employer should be able to state which duty of the job requires that attribute and on what evidence. Published research documents substantial word error rate disparities in automated speech recognition for speakers whose speech falls outside the training distribution, including racialized speech patterns and atypical or disordered speech, which is why the University of Illinois established a dedicated Speech Accessibility Project in 2022 to collect such data. Where a candidate signs through an interpreter, the transcript being scored is the interpreter's rendering rather than the candidate's utterance, and the interpreter's lexical choices and processing lag are attributed to the candidate. The system produces a number. It is not measuring the candidate.

The state layer and its timetable

Federal retrenchment does not clear the field, and the state schedule matters for planning. Illinois House Bill 3773, signed in August 2024, amends the Illinois Human Rights Act to prohibit discriminatory use of artificial intelligence across the employment lifecycle and to require notice to applicants and employees when such systems are used, effective January 1, 2026. Colorado's Senate Bill 24-205, signed in May 2024, imposes duties on developers and deployers of high-risk systems used in consequential decisions including employment,

with an effective date of February 1, 2026, though the governor invited legislative revision at signing and amendment before that date should be considered likely. Virginia's House Bill 2094, which would have established a comparable framework, was vetoed in March 2025. New York City's Local Law 144 continues in force.

The direction of the state layer is toward notice and disclosure rather than substantive duties of care. Employers should plan on the assumption that the obligations they will actually face are to say what they are doing, and that the obligation to be able to prove what they did will arrive through litigation rather than through a regulator.

What to do now

The following can be completed by an existing human resources function and does not require external assistance.

1. Inventory every system that touches a selection decision, including résumé parsing, ranking, knockout questions, chat-based prescreening, and automated video or speech assessment. Include features enabled by default in an applicant tracking system, which are the most commonly overlooked.
2. For each system, state in writing what it measures and which documented job requirement that construct corresponds to. Where no answer exists, the system is scoring an unvalidated criterion.
3. Confirm that candidates are told an alternative or modified assessment is available before the assessment begins, not after, and that the notice is on the assessment itself rather than buried in a careers page.
4. Retrieve every accommodation request associated with an automated assessment in the past twenty-four months. Record what was requested, who decided, what was provided, and what the outcome was. This file will be produced in any charge and should be reviewed before someone else reviews it.

5. Ask the vendor, in writing, for validation evidence for each scored construct and for accuracy data disaggregated by speech and communication characteristics. Retain the answer, including a refusal, which is itself useful evidence of the state of the employer's knowledge.
6. Confirm that scores and the configuration in force at the time of each decision are retained for the applicable preservation period. Scores computed to produce a ranking and then discarded cannot be explained later.
7. Treat any bias audit result as evidence about race and sex categories only, and do not represent it internally or externally as evidence of ADA compliance.

THE THREE QUESTIONS

Was the assessment accessible to this candidate. Was an accommodation requested, and what was the response. Can the employer explain what the assessment measured and why that construct is job-related.

Conclusion

Artificial intelligence can materially expand access to work for disabled people. Real-time captioning, speech-to-text, and adaptive interfaces have already done so, and the OECD's 2023 assessment of AI and disability in the labour market documents genuine gains. Those gains come from assistive applications, where the technology serves the person. They do not transfer automatically to evaluative applications, where the technology measures the person, and the current tendency to argue from the first to the second is the central analytic error in this field.

The withdrawal of federal guidance in January 2025 changed what employers are told. It did not change what employers owe. The organizations that will be able to defend their hiring

decisions are not the ones that bought the most sophisticated tool or published the most reassuring audit. They are the ones that can say what their systems measure, why it matters to the job, what happened when a candidate asked for something different, and who decided.

References

- American Civil Liberties Union, ACLU of Colorado, Public Justice, and Eisenberg & Baum LLP. Complaint filed with the Colorado Civil Rights Division and the U.S. Equal Employment Opportunity Commission, March 19, 2025.
- Americans with Disabilities Act of 1990, Title I, 42 U.S.C. § 12112.
- Clifford Chance. *Inclusive AI for People with Disabilities: Key Considerations*. 2024.
- Colorado Senate Bill 24-205, Concerning Consumer Protections in Interactions with Artificial Intelligence Systems, signed May 17, 2024.
- Executive Order 14179, Removing Barriers to American Leadership in Artificial Intelligence, January 23, 2025.
- Illinois House Bill 3773, amending the Illinois Human Rights Act, signed August 9, 2024.
- New York City Local Law 144 of 2021, Automated Employment Decision Tools.
- OECD. *Using AI to Support People with Disabilities in the Labour Market*. 2023.
- U.S. Equal Employment Opportunity Commission. *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees*, EEOC-NVTA-2022-2, May 12, 2022. Removed from the agency website January 27, 2025.
- Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. Part 1607; recordkeeping requirements, 29 C.F.R. Part 1602.