

Blog Post

# AI in Hiring: Emerging Legal Developments and Compliance Guidance for 2026

November 20, 2025

By [Reeya Khurana](#)

AI isn't just on the horizon—it's already screening millions of resumes, scoring video interviews, and ranking candidates in HR systems across America. In 2024 alone, AI-powered hiring tools processed over 30 million applications while triggering hundreds of discrimination complaints. As these tools become more prevalent, lawmakers, regulators, and attorneys are responding rapidly. The result is a legal landscape evolving faster than most compliance teams can track. For employers, staying informed isn't optional—it's essential. Here's what to expect in the year ahead.

[State Law Showdown: What's on the Books and What's Coming Next?](#)

[New York City Local Law 144 \(Effective July 2023\)](#)

New York City set the tone for the nation with a law requiring annual, independent bias audits for any automated employment decision tools (AEDTs) used in hiring or promotion.

Employers must post audit summaries online, notify candidates and employees at least 10 business days before using an AEDT, and offer an alternative

---

## Related People

[Reeya Khurana](#)

---

## Related Work

[Employment Administrative Claims Defense](#)

[Employment Training and Compliance Labor and Employment](#)

---

## Related Offices

[New York](#)

---

## HR Defense

[Akerman Perspectives on the Latest Developments in Labor and Employment Law](#)

[Visit this Akerman blog](#)

selection process if requested. Fines escalate rapidly: \$500–\$1,500 per violation, multiplied by each day of non-compliance and each affected applicant—potentially reaching millions for systematic violations.

Notably, NYC Department of Consumer and Worker Protection guidance confirms that the law applies even when humans make final decisions based on AI rankings or scores. The human’s role doesn’t eliminate the need for compliance if AI tools are influencing the decision-making process.

### California Civil Rights Council Regulations (Effective October 1, 2025)

California’s rules are the most detailed yet. The regulations state that it is unlawful to use any “automated-decision system” (ADS) that discriminates against applicants or employees based on protected traits when making hiring or personnel decisions.

Accordingly, any ADS used in employment must have meaningful human oversight, with someone trained and empowered to override the AI. Employers must proactively test for bias, keep detailed records for at least four years, and provide reasonable accommodations or alternative assessments if an ADS could disadvantage people based on protected traits.

The regulations also clarify that any ADS that elicits information about an applicant or employee’s disability may constitute an unlawful medical inquiry.

Notably, the regulations make clear that vendors and software providers can be held liable under traditional agency principles when they exercise control over employment decisions or act on behalf of the employer in recruitment or screening. This expands potential liability beyond just the hiring employer.

## Texas Responsible Artificial Intelligence Governance Act (TRAIGA) (Effective January 1, 2026)

In contrast, Texas has taken a minimalist approach. TRAIGA establishes a general framework for AI development, government transparency, and consumers rights; however, it imposes fewer restrictions on private sector employers.

Fundamentally, TRAIGA bans intentional discrimination via AI but rejects disparate impact as a standalone basis for liability. Therefore, merely showing that an AI system negatively impacts a protected class would not, by itself, establish a violation—a significant philosophical departure from California’s approach.

The state attorney general has exclusive enforcement power, and employers get a notice and a 60-day cure period before penalties kick in. Fines can range from \$12,000 for violations the court determines are curable, up to \$200,000 for violations the court deems incurable, and up to \$40,000 per day for continuing violations.

## Illinois House Bill 3773 (Effective January 1, 2026)

Illinois is addressing concerns about discriminatory practices stemming from misuse of AI. Under the new law, employers can’t use AI in ways that result in bias against protected classes under the Illinois Human Rights Act, whether intentional or not, and must notify employees and candidates when AI is used in employment decisions.

Using ZIP codes as a proxy for protected characteristics is also banned—a recognition that discriminatory outcomes can occur even when protected characteristics aren’t directly considered.

While the law does not allow private lawsuits, individuals can file complaints with the Illinois

Department of Human Rights.

## Colorado Artificial Intelligence Act (SB 24-205) (Effective June 30, 2026)

Colorado’s law is among the most comprehensive passed by the states. It regulates the use of “high-risk” AI systems—any AI that makes or influences significant employment decisions like hiring, firing, or promotion—to ensure that high-impact hiring tools are used in a fair, transparent, and legally compliant manner.

Violations constitute an unfair trade practice under Colorado’s Consumer Protection Act. Under the law, vendors and employers must express foreseeable uses and risks to applicants and employees, complete an annual impact assessment, and maintain transparency when individuals interact with an AI system or when the system makes an adverse decision.

A complainant must notify the attorney general within 90 days if discrimination is discovered. If an adverse decision is made, documentation and an appeal process are required.

### Common Threads in State AI Legislation: Tackling Bias and the “Black Box”

While these five states’ approaches differ in their details, they reveal shared concerns that every employer should understand. A clear pattern is emerging: lawmakers are focused on two major interconnected concerns that make AI-driven hiring particularly challenging to regulate.

#### **The Transparency Challenge**

The first is the well-recognized challenge of AI transparency—often referred to as the “black box” issue. AI systems can be so complex that even their developers struggle to explain how specific decisions are made. This lack of transparency makes

it difficult for employers to understand, audit, or defend the reasoning behind automated hiring or promotion outcomes.

## **How AI Perpetuates Bias**

The risk of perpetuating or amplifying bias is another concern, and understanding how this happens is crucial. AI tools are typically trained on large datasets of historical hiring decisions. If those datasets reflect past inequities—such as periods when women were rarely hired for technical roles or older workers were systematically passed over for promotions—the AI learns to replicate those patterns.

The AI doesn't understand context or history; it simply identifies correlations in the data. When it observes that, historically, successful candidates for engineering roles were predominantly male, it may incorrectly learn that male gender correlates with engineering success. The system isn't explicitly programmed to discriminate—it's learning discriminatory patterns embedded in historical data.

Even more problematically, bias persists even when explicit references to protected characteristics are removed from the system. The AI identifies proxy variables—seemingly neutral factors that correlate with protected traits. For instance, graduation years can serve as proxies for age, certain ZIP codes correlate with race, and gaps in employment history may disproportionately affect women who took parental leave. The AI uses these proxies to make predictions, inadvertently perpetuating discrimination.

## **Why This Matters Legally**

Because employment decisions are governed by a robust framework of anti-discrimination laws—including Title VII of the Civil Rights Act, the Americans with Disabilities Act, and the Age

Discrimination in Employment Act—these risks are not just theoretical.

When a “black box” algorithm relies on biased or proxy-laden data, employers may unknowingly run afoul of both federal and state anti-discrimination statutes, exposing themselves to significant legal liability. The opacity of the AI system makes it nearly impossible for employers to verify that their hiring processes are legally compliant.

## Courtroom Chaos: Legal Scrutiny for AI-Driven Employment Decisions

The legal spotlight is now shining on a growing list of high-stakes cases, and the trends emerging from such litigation should concern every employer using AI hiring tools.

### **Vendor Liability: A New Frontier**

Recent federal litigation has established that HR technology vendors themselves can be held liable for discriminatory outcomes, not just the employers who use their tools. In cases involving allegations that AI-powered tools contributed to disproportionate screening outcomes—including age-based rejection of qualified applicants—courts have recognized that technology vendors can face liability under traditional agency principles or as entities exercising control over employment decisions.

In at least one significant case, the court ordered a vendor to produce a list of all customers using its AI features over a multi-year period, dramatically raising the stakes for both vendors and employers. The court’s willingness to pursue vendor liability and to expose the vendor’s entire customer base to potential scrutiny represents a watershed moment in AI employment litigation.

### **EEOC Enforcement Actions**

The Equal Employment Opportunity Commission has reached settlements with companies after alleging that their AI-powered application processes automatically screened out applicants based on protected characteristics such as age. These enforcement actions signal that federal agencies are actively investigating AI hiring systems and won't hesitate to take action when they identify discriminatory patterns.

## Accessibility and Bias Concerns

Additional complaints have raised concerns about AI-powered video interview tools that may be inaccessible to candidates with disabilities—such as deaf applicants—or that produce biased results for individuals from diverse linguistic or cultural backgrounds. These cases highlight that AI bias isn't limited to traditional demographic categories; it can manifest in unexpected ways that disadvantage candidates based on speech patterns, accents, or communication styles.

## What These Cases Mean for Employers

These cases are just the tip of the iceberg, and together they signal a new era of legal scrutiny for AI-driven employment decisions. Several key principles are emerging:

- **Transparency and Explainability Are Critical.** Employers should expect to be asked for detailed documentation on how AI tools are used, how decisions are made, and what steps have been taken to monitor for bias. Regulators and plaintiffs alike are seeking to uncover the “black box” of AI decision-making, and employers who can't explain or justify their systems may find themselves at a disadvantage in court.
- **Using a Vendor's Tool Doesn't Shield You.** Courts are signaling that both the employer and the technology provider can be held responsible if AI-driven decisions result in discrimination. The

“we’re just using a vendor’s tool” defense is rapidly eroding.

- **Documentation Is Your Best Defense.** The employers who fare best in this litigation have maintained detailed records of bias testing, have documentation showing meaningful human oversight, and can demonstrate that they’ve proactively addressed identified problems.

## Navigating the Multi-State Compliance Challenge

For employers operating across multiple states, the patchwork of regulations creates significant compliance challenges. A hiring system that meets Texas’ minimal requirements may fall far short of California’s comprehensive standards.

Consider a national retailer using AI tools for high-volume hiring across all 50 states: they must implement California-level bias auditing, testing, and documentation to remain compliant in that state. But, those same robust procedures become a competitive advantage when defending against claims in other jurisdictions.

The practical approach? Design your AI governance program to meet the most stringent state requirements—currently California’s—and you’ll achieve compliance nationwide while building the strongest possible defense against litigation.

## Five Steps to Navigate AI Compliance

- **Map and Monitor Your AI:** Maintain a current, comprehensive inventory of every system—internal or vendor-supplied—that scores, ranks, filters, or evaluates candidates or employees. This includes obvious tools like resume screening software, but also less obvious systems like automated interview scheduling that might use algorithms to rank candidates. Regularly review and update this inventory as your tech stack evolves. Many employers are surprised to

discover how many AI touchpoints exist in their hiring process when they conduct this exercise.

- **Establish Robust AI Governance:** Develop a formal governance program that assigns clear roles and responsibilities for AI oversight, sets policies for evaluating and deploying AI tools, and ensures leadership buy-in for compliance. This isn't just a legal or HR function—it requires coordination across IT, procurement, legal, HR, and business units. Someone at the senior level should own AI governance and have the authority to pause or discontinue problematic tools.
- **Audit for Bias and Compliance:** Conduct regular, documented audits of all AI tools for bias and disparate impact, using both internal and, where possible, third-party experts. Be prepared to show your work if challenged. Effective audits measure selection rates across protected groups, test for disparate impact using statistical significance tests, and probe for proxy variable effects. Document not just what you found, but what you did in response to concerning findings.
- **Strengthen Vendor Contracts and Oversight:** Require vendors to provide transparency into their AI systems, share audit results, commit to ongoing bias testing, and accept contractual liability for discriminatory outcomes. Don't rely on "black box" assurances or vendor claims that their systems are "bias-free." Insist on seeing the methodology, the testing data, and the audit results. Include provisions requiring vendors to notify you immediately if they discover bias issues, and establish clear procedures for addressing problems.
- **Communicate and Train:** Clearly notify candidates and employees when AI is used in employment decisions, explaining how the tools work and how decisions are made. Train HR and management teams on AI bias risks, legal requirements, and how to respond to regulatory

inquiries or litigation. Ensure that the humans providing “oversight” of AI systems actually understand what they’re overseeing and have the authority and training to override AI recommendations when appropriate.

## The Benefits of Thoughtful AI Implementation

These compliance requirements shouldn’t obscure an important reality: AI tools, when thoughtfully implemented with appropriate safeguards, offer genuine benefits.

They can help organizations efficiently manage large applicant pools that would be impossible to review manually, potentially opening opportunities for candidates who might otherwise be overlooked in purely human screening. They can reduce certain forms of human bias by ensuring consistent evaluation criteria. And they can improve the candidate experience by providing faster responses and more efficient scheduling.

The goal of these regulations isn’t to eliminate AI from hiring—it’s to ensure that when these powerful tools are used, they’re deployed responsibly, transparently, and in ways that expand opportunity rather than perpetuate historical discrimination.

## The Path Forward: Managing AI Hiring Tools in a Changing Landscape

The age of unregulated AI in employment is over. As we move into 2026, employers face a patchwork of state, local, and (potentially) federal requirements designed to make AI-driven employment decisions fair, transparent, and accountable.

Lawsuits and regulatory scrutiny are on the rise, and the risks of non-compliance are real—both in terms of direct penalties and reputational harm. But, employers who take proactive steps now can harness genuine efficiency and consistency benefits of AI without becoming the next cautionary tale.

Expect federal legislation to emerge by late 2026 or early 2027, likely attempting to harmonize these patchwork requirements. Based on current momentum, federal standards will probably follow California's comprehensive model more closely than Texas' minimal approach, making early investment in robust compliance programs a sound strategic choice.

The employers who will thrive in this environment are those who view AI compliance not as a burden but as a competitive advantage—demonstrating to candidates, regulators, and the public that they're using powerful technology responsibly and ethically.

### Take Action Now

If you're using or considering AI hiring tools, the time to act is before you receive a complaint or a regulatory inquiry. Akerman's Labor & Employment team has deep experience helping clients navigate these complex requirements across multiple jurisdictions. If you have any questions or want to discuss how these requirements apply to your specific situation, please reach out to your Akerman Labor & Employment attorney. Because when it comes to AI, it pays to have a real human—with real legal expertise—in your corner.

---

This information is intended to inform firm clients and friends about legal developments, including recent decisions of various courts and administrative bodies. Nothing in this Practice Update should be construed as legal advice or a legal opinion, and readers should not act upon the information contained in this Practice Update without seeking the advice of legal counsel. Prior results do not guarantee a similar outcome.