HR Wellness Foundation Series | Monthly Blog

AI in Hiring: What Startups Should Know Before Using It

Written by: Sophia Fulton | Chameleon Strategies Consulting

There is a version of AI-assisted hiring that sounds almost too good to be true: faster screening, reduced administrative burden, smarter candidate matching, and more time for the human conversations that actually matter. And in some ways, that promise is real. AI tools have genuinely changed what is possible in talent acquisition, particularly for small teams stretched thin across multiple priorities.

But there is another version of this story that is not talked about nearly enough... The one involving discrimination lawsuits, regulatory penalties, recruiting fraud, and legal exposure that most early-stage companies never see coming until they are already in it.

This month, we are talking about what life sciences startups need to understand before they bring AI into their hiring process. Because the decision to adopt these tools is much more than a technology decision. It is a legal and ethical one.

The Regulatory Landscape Has Already Shifted

If you have not been tracking state-level AI hiring legislation, now is the time to start. The regulatory environment has moved faster than most compliance teams, let alone early-stage HR functions, can keep up with.

New York City's Local Law 144, in effect since 2023, prohibits employers from using any automated employment decision tool unless it has undergone an independent bias audit within the past year, with audit results published publicly on the employer's careers page, and candidates notified at least ten business days in advance of its use (Maurer, 2023). Non-compliant organizations face fines of up to $500 for a first violation and $1,500 for every subsequent day the unaudited tool remains in use. Illinois amended its Human Rights Act effective January 1, 2026, to expressly prohibit the use of AI that produces discriminatory effects in recruiting and hiring, and requires employers to notify candidates when AI will be used in the process, a requirement that extends to any employer with even a single employee in Illinois (McRee et al., 2024). California's Civil Rights Council regulations, effective October 1, 2025, clarify that existing anti-discrimination law applies when employers use automated decision systems, including when those systems produce disparate impact on protected characteristics, regardless of intent (Khuarana, 2025). Colorado's original Artificial Intelligence Act classified employment-related AI as 'high-risk,' prohibited its use without documented safeguards, and required annual impact assessments and candidate notifications (Khuarana, 2025). But on May

14, 2026, Governor Polis signed SB 26-189, replacing that framework entirely with a significantly scaled-back law focused on transparency and disclosure rather than bias audits or risk assessments — effective January 1, 2027 (Tyman et al., 2026). Colorado's rapid reversal is a reminder that this regulatory landscape is not just evolving; it is actively being rewritten, sometimes within months.

The pattern across all of these laws is consistent and clear: transparency, human oversight, and accountability must be at the heart of AI use in the hiring process. And critically, employers recruiting remote workers may be subject to these laws even if they are not physically located in the regulating jurisdiction (Galkin, 2026). For a life sciences startup hiring across state lines, which most do, this creates a compliance patchwork that demands attention even at ten employees.

Federal legislation is likely next, so organizations that build compliance practices now will not have to scramble later.

The Algorithm Is Not Neutral, and You Are Still Liable

One of the most consequential misunderstandings about AI hiring tools is that because a third-party vendor built the tool, the employer is insulated from liability for how it performs.

This is absolutely not the case.

Anti-discrimination laws apply regardless of whether bias is generated by a human or an algorithm, a principle the EEOC established explicitly when it launched its AI fairness initiative in 2021 (U.S. EEOC, 2021). If an algorithm screens out candidates in ways that disproportionately affect a protected class, by gender, race, age, or disability status, liability attaches to the employer regardless of whether the tool was built internally or procured from a vendor (Galkin, 2026). The ongoing Workday litigation, in which a nationwide class action was certified on behalf of job seekers over forty alleging algorithmic age discrimination, is a real-world example of what this exposure looks like at scale (Rigney et al., 2026).

The implications for startups are practical and immediate. Before adopting any AI hiring tool, founders and HR leaders should understand:

- **What data the tool uses to make decisions.** AI systems trained on historical hiring data can encode and replicate the biases of past decisions. If your industry historically underrepresented women in technical roles, a tool trained on that data may systematically deprioritize female candidates.
- **Whether the vendor conducts bias audits.** Reputable vendors should be able to demonstrate that their tools have been tested for disparate impact across protected classes and provide documentation of those audits.

- **Whether meaningful human oversight exists.** Nearly every state law and regulatory framework now requires that a trained human person be empowered to review and override AI recommendations. A fully automated adverse decision, one that rejects a candidate without any additional human review, creates significantly higher legal exposure.
- **What your data retention obligations are.** For example, California's automated decision system regulations require employers to maintain employment records, including data used by or generated by automated decision systems, for a minimum of four years. Know what your vendor retains, where, and for how long, and ensure your own recordkeeping practices align with the states where you hire (Hoopai & Ascar, 2026).

The compliance landscape has caught up with the technology. What once felt like a gray area is now clearly governed, and adopting a tool without understanding how it works carries real legal and ethical consequences.

The Recruiting Scam Problem No One Is Talking About

AI has made the hiring process more efficient for employers. But it has also made it significantly easier for bad actors to exploit candidates, and to do so using your company's name.

Recruiting fraud is growing at a rate that should concern every organization actively hiring. Between May and July of 2025 alone, job scams increased by more than 1,000%, according to a McAfee report (Schuman, 2026). The FTC recorded over 25,000 reports of business and job opportunity scams in the last quarter of 2025 alone, resulting in more than $150 million in losses and a median individual loss of approximately $2,000 (Taylor, 2026). And because fewer than 10% of fraud victims ever report to a federal agency, the true scale of this problem is almost certainly far higher.

Here is how these scams typically work: a candidate receives outreach from someone claiming to be a recruiter for a company they recognize, often a legitimate life sciences organization actively posting roles. The message looks professional, the "recruiter" may have a polished LinkedIn profile, and an offer follows quickly. Then comes the request: banking information for direct deposit, a Social Security number for onboarding paperwork, or a fee for a background check or training. But in reality, there is no job, and there never was.

AI has made this problem significantly harder to detect. Today's recruiting scams use large language models to scrape a candidate's LinkedIn profile and generate personalized outreach that mirrors their professional history and career trajectory, making the approach feel eerily tailored and credible (Schuman, 2026). Deepfake video interviews, in which fraudsters use real-time face-swap technology to impersonate company executives during Zoom calls, have also been documented and are expected to continue to grow.

For startups, this is both a candidate protection issue and a serious brand risk. When someone is defrauded using your company's name, the damage extends beyond the individual victim. It erodes trust in your employer brand, creates confusion in your candidate pipeline, and may require active intervention to manage.

What your organization should do:

- **Add explicit language to all job postings and your careers page** stating that your recruiting is conducted internally (or through named, authorized partners only), that you will never request banking information, Social Security numbers, or fees during the application process, and that candidates who receive suspicious outreach claiming to be from your organization should report it directly to a designated HR contact.
- **Train your recruiting team** to be aware that your company name may be used fraudulently and establish a clear internal process for responding when candidates flag suspicious contact.
- **Audit your digital presence.** A clear, current careers page with named HR contacts makes it easier for candidates to verify what legitimate outreach looks like and much harder for scammers to impersonate it convincingly.
- **Report confirmed scams** to the FTC at ftc.gov/complaint and notify the platform through which the fraud occurred.

A simple disclosure statement added to your job postings costs nothing. For a candidate who almost gave their bank account information to a fraudster using your company's name, it could make all the difference.

What Responsible AI-Assisted Hiring Actually Looks Like

None of this means AI has no place in early-stage hiring. Used thoughtfully, these tools can genuinely support efficiency and consistency — two things early-stage teams desperately need. The question is not whether to use AI, but how to use it safely.

- **Start with transparency.** Tell candidates when AI is being used in your process. In several states this is now legally required, but beyond compliance, it is simply the right approach. Candidates deserve to know how their applications are being evaluated.
- **Keep human people in the loop for every consequential decision.** AI can surface candidates, flag patterns, and reduce manual screening time, but it should not be the final word on who gets an interview or an offer. Every adverse decision, every rejection, should involve human review.

- **Vet your vendors before you sign.** Ask directly about bias testing, data retention, audit history, and what happens if their tool produces a discriminatory outcome. A vendor that cannot answer these questions clearly is a vendor that has not thought carefully enough about these risks.
- **Document your process.** The regulatory environment is moving toward accountability, and documentation is your protection. Record how AI tools are used, what decisions they inform, and what human oversight looks like at each stage.
- **Review regularly.** A tool that performed well at hire may not perform the same way as your organization scales, roles evolve, or the underlying model is updated. Build review into your process, not just your onboarding.

The Bottom Line

The promise of AI in hiring is real, but so is the risk. For early-stage life sciences companies, the combination of a rapidly shifting regulatory landscape, meaningful legal liability, and a surge in recruiting fraud means that adopting these tools without a clear-eyed understanding of what they require is genuinely dangerous.

The organizations that will navigate this moment most successfully are not the ones that avoid AI out of fear or adopt it without guardrails out of enthusiasm. They are the ones that treat it the same way they would any other consequential HR practice: with intention, documentation, human oversight, and a commitment to protecting both the company and the people it is trying to hire.

A question to sit with this month: If a candidate asked you today to explain exactly how AI is, or isn't, being used in your hiring process, would you be able to answer clearly and confidently?

If the answer is uncertain, that is exactly where the work begins.

References

Galkin, W. S. (2026, January 16). *2026 overview of AI use in employment decisions.* GalkinLaw LLC + AI Governance Team via Lexology. https://www.lexology.com/library/detail.aspx?g=bb0a51a8-4a1f-4592-83a2-3de69f22d075

Hoopai, L., & Ascar, C. (2026, June 3). *How Illinois, Colorado, California, and Texas approach state AI hiring laws.* DISA. https://disa.com/news/ai-hiring-laws-by-state/

Khuarana, R. (2025, November 20). *AI in hiring: Emerging legal developments and compliance guidance for 2026.* Akerman LLP / HR Defense Blog. https://www.hrdefenseblog.com/2025/11/ai-in-hiring-emerging-legal-developments-and-compliance-guidance-for-2026/

McRee, E. B., Andrianifahanana, E., Butler, W. C., Linas, J. M., Ling, J. C., Paez, M. F., & Schulman, E. R. (2024, October). *Illinois becomes second state to pass broad legislation on the use of AI in employment decisions.* Jones Day. https://www.jonesday.com/en/insights/2024/10/illinois-becomes-second-state-to-pass-broad-legislation-on-the-use-of-ai-in-employment-decisions

Maurer, R. (2023, April 16). *NYC outlines compliance requirements for AI bias law.* SHRM. https://www.shrm.org/topics-tools/news/talent-acquisition/nyc-outlines-compliance-requirements-ai-bias-law

Rigney, E. L., Sparhawk, I. F., Chiu, S. G., Cogan, R. J., Diaz, R. E., Healy, M. C., Lagunas, Y., Meredith, M. W., & Velasquez, C. M. (2026, February 3). *Navigating the AI employment landscape in 2026: Considerations and best practices for employers.* K&L Gates. https://www.klgates.com/Navigating-the-AI-Employment-Landscape-in-2026-Considerations-and-Best-Practices-for-Employers-2-2-2026

Schuman, N. (2026, April 10). *Recruitment scams survival guide: Outsmarting deceptive AI.* Lloyd Staffing. https://www.lloydstaffing.com/recruitment-scams-survival-guide-outsmarting-deceptive-ai/

Taylor, B. (2026, June 2). *The job offer may be fake: 4 recruitment scams experts say are on the rise.* Click2Houston. https://www.click2houston.com/news/local/2026/06/02/the-job-offer-may-be-fake:-4-recruitment-scams-experts-say-are-on-the-rise/

Tyman, A., Vele, J. R., & Laudenbach, C. J. (2026, May 22). *Colorado enacts artificial intelligence replacement law.* Seyfarth Shaw LLP. https://www.seyfarth.com/news-insights/colorado-enacts-artificial-intelligence-replacement-law.html

U.S. Equal Employment Opportunity Commission. (2021, October 28). *EEOC launches initiative on artificial intelligence and algorithmic fairness.* https://www.eeoc.gov/newsroom/eeoc-launches-initiative-artificial-intelligence-and-algorithmic-fairness