When the Algorithm Recommends Termination: Employer Liability for AI in Hiring, Discipline, and Performance Management

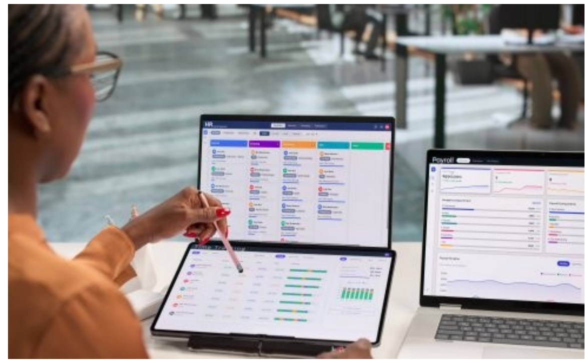

August 27, 2026
By Sarah Goodman

Employers have moved quickly to bring artificial intelligence into parts of the employment relationship that used to depend entirely on human judgment. Applicant tracking systems now score resumes before a recruiter sees them. Scheduling and productivity platforms flag employees as underperforming based on keystroke counts, call times, or delivery windows. Performance management tools generate draft write-ups, and in some organizations, recommend whether an employee should be coached, placed on a performance improvement plan, or terminated. The efficiency case for these tools is obvious. The legal exposure they create is less obvious, and it is growing.

The core problem is not that AI is involved in employment decisions. It is that AI outputs are increasingly being treated as conclusions rather than inputs, and that shift changes how those decisions look in a deposition, an EEOC position statement, or a jury instruction.

The Employer Cannot Delegate the Decision

Title VII, the ADA, the ADEA, and their state and local counterparts all impose liability on the employer, not on the software the employer purchased. An employer cannot defend a discrimination claim by pointing to a vendor's algorithm and arguing that a machine, not a person, made the call. Regulators have already made this point explicit. The EEOC has stated that employers remain responsible for adverse impact caused by algorithmic decision-making tools even when a third-party vendor built and maintains the tool. Several state laws, including New York City's Local Law 144 governing automated employment decision tools, impose independent audit and notice obligations directly on the employer using the tool.

That means the familiar advice to document a legitimate, nondiscriminatory reason for an adverse action now has a second layer. It is not enough that AI flagged the employee. The employer must show that a human reviewed the flag, understood why it was generated, and exercised independent judgment before acting on it. A termination file that says only "system-generated performance score of 2.1, employee terminated" is a much harder file to defend than one that documents what a manager actually observed and considered.

Disparate Impact Hides Inside the Model

Disparate treatment claims require some evidence of intent, but disparate impact claims do not, and AI-driven employment tools are a natural fit for disparate impact theories. A scoring model trained on historical performance or attrition data can quietly reproduce whatever bias existed in that history. A resume screening tool can learn to penalize employment gaps, certain schools, or language patterns that correlate with protected characteristics even though the model was never told to consider race, sex, age, or disability directly.

The practical exposure here is twofold. First, if a plaintiff's counsel obtains statistical evidence that an AI tool selects or rejects candidates or employees at meaningfully different rates across protected groups, the employer will need validation data showing the tool is job related and consistent with business necessity, the same standard that has applied to any selection device since Griggs v. Duke Power Co. Vendors rarely provide this validation data unprompted, and employers frequently discover during litigation that they never asked for it. Second, an employer that never tested its own tool for adverse impact will have a difficult time arguing it acted reasonably, even where no discriminatory intent existed. Ignorance of how the tool works is not a defense; in a disparate impact case, it can be the plaintiff's best evidence that no one was minding the store.

Discovery Now Reaches Further Than the Personnel File

AI adoption expands what is discoverable in an employment case well beyond the traditional personnel file. Prompts entered by HR or supervisors, model outputs and confidence scores, version histories showing when a scoring model was retrained, and internal communications about why an alert was or was not acted on are all now fair game. Plaintiffs' counsel is increasingly requesting this material specifically, because it can show not just that an adverse outcome occurred, but what the company knew and when.

This creates a preservation problem many employers have not yet solved. Some platforms overwrite scoring history as new data comes in or retain outputs only briefly by default. If a company has no policy governing retention of AI-generated employment data, it may find itself unable to produce records that plaintiffs assume exist, inviting a spoliation argument, or it may find that the only surviving record is an unfavorable one that a human reviewer never actually relied upon but that now looks, in hindsight, like the smoking gun.

Practical Governance Closes the Gap

None of this counsels against using AI in employment decisions. It counsels against using it without a governance structure built for the way these tools will actually be examined after the fact. At a minimum, employers should be able to show that a person with real authority to disagree reviewed any AI-generated recommendation before it became a personnel action, that the tool has been validated or at least tested for disparate impact on a periodic basis, that managers are trained to document their own independent reasoning rather than simply citing the tool's output, and that retention practices for AI-generated data are deliberate rather than accidental. Where a jurisdiction imposes specific notice, bias audit, or disclosure obligations on automated employment decision tools, those requirements need to be built into the rollout, not addressed after a candidate or employee complains.

The employers best positioned when one of these tools produces a bad outcome are the ones who can show, with contemporaneous documentation, that a human being was accountable for the decision the whole way through. The tool can inform that judgment. It cannot substitute for it, and the law has no intention of letting it try.