

The Algorithm Ate My Residency (And Other Stories for Employers)

Mark A. Konkel

May 11, 2026

Employers using AI-powered tools to screen, rank, or normalize job applicants face significant and growing litigation risk—particularly when they cannot explain how algorithmic decisions are made. A recent incident involving Thalamus Cortex, an AI platform used in the medical residency matching process, illustrates how opaque algorithms can systematically disadvantage qualified candidates in ways no human reviewer would replicate. Companies that deploy these tools without understanding what happens “under the hood” should expect data access requests, discrimination claims, and class-action exposure.

We don't like to overreact here at [Labor Days](#). We do believe that AI is transformative in the world, the workplace, and employment. We do not believe that we are all zombies who died when OpenAI launched and just don't know it yet.

But a recent report floating around the medical residency world—specifically involving a tool called Thalamus Cortex—has now made old fears seem quaint. It takes the “AI-gone-wrong” riskometer from “omg, that guy has six fingers” to “we are accidentally destroying the next generation of brain surgeons.” That's a pretty big liability gap.

A WIRED report (which, for the record, is very much a real thing, backed by a study in *The Laryngoscope*) describes a nightmare scenario for any company or employer considering applications. (Which is . . . all of them.) Physicians at a national meeting noticed that applicant grades in the Cortex system were literally shifting before their eyes. [Cue horror track: the humans watch AI turn on them in real time.] One minute you're a straight-A superstar; the next, the AI “normalizer” has said you're a C+ on a good day.

The CEO of the company told WIRED it was basically a UI glitch—the digital equivalent of the screen getting “stuck” because the user clicked too fast. If you're tracking: you've spent hundreds of thousands of dollars on a medical degree; you've worked harder until now than most people work in a lifetime; and your career—remember, where you land now is the starting point for all potential future roads—has been gutted because the system was slow, and you chose to do the thing, the one thing, that no person must ever do: you clicked twice.

Saying this happened because the “user clicked too fast” is the tech version of “the dog ate my homework,” except the homework is someone's life.

Enter Chad Markey, the hero of our story. Markey overcame genuinely extraordinary personal odds to graduate from an elite medical school with glowing recommendations—a sterling candidate. When similar peers got interviews and he faced nothing but rejections, he suspected the AI visible in

the application process was penalizing him for mentioning a leave of absence for “personal reasons.” He had listed those periods as leave for ‘personal reasons’; in reality, he was obtaining medical treatment for an extremely serious, life-threatening illness, and included the more generalized description of the leave. He asked what happened. He didn’t get answers. [Cue dramatic track.].

So he got even. He went full Mr. Robot.

He used Python to create a synthetic dataset of 6,000 applicants and reverse-engineered the company’s own patents using Claude Code and GitHub Copilot. What he did is beyond impressive, but what it revealed is fairly obvious from all prior human experience: language matters.

It matters a lot, apparently. As it turns out, the AI’s “sentiment analysis” reportedly favored “medically accurate” descriptions over “personal reasons” by a massive margin. That was why. So a general coding instruction to prefer more specific or accurate information over less specific or potentially less accurate information—sensible at the general level—worked its way through algorithmically to apply that criterion to a person whose medical information is personal to him, and whose personal identity might have been organized around working through challenges rather than announcing them. And then rejected him.

A person would not have done that.

The company denies that their AI actually “scored” or “ranked” Markey’s specific cycle. That isn’t entirely persuasive, apparently, because they do use AI for what they call “grade normalization”—which is related to the idea of leveling the playing field or “comparing apples to apples.”

But they learned the hard algorithmic lesson: putting inputs into an algorithm hands those inputs over to a new process, an exponential process, that we don’t watch or really understand. So we aren’t generally aware of just how much “thinking” is going on until it goes wrong: a hand with six fingers or an unemployed Josh Markey. If you’re Chad Markey, AI might even make the exact opposite choice you would make based on the same information. In Markey’s case, the record of his perseverance in a demanding profession was exactly the thing that Cortex noticed to disqualify him.

Why Should Employers Care About AI Hiring Tools?

The usual, self-explanatory word: litigation.

When employers use AI, they are using a process whose moment-to-moment functioning, as used by real people, is not known. Here, let’s be intellectually honest with the word that Cortex used, one that is used a lot to try to pinpoint exactly what we don’t know about AI: “opaque.” Obviously, AI is human-designed and human-made, so it wouldn’t be true to say that we literally don’t know “how” AI works. But AI behaves independently, even if we made it. That’s the point. It’s a complex algorithm running in the background performing millions of operations you never see and synthesizing all that into outputs that feel human, or seem to have and provide exactly the information you asked for (assuming it’s not hallucinated).

In our experience, claimants sue companies when they lack power or information, or both. Employers have always had more power and information. Employers must acknowledge the obvious: humans know less about the systems they’re working within when they use AI at scale. So AI will legitimately weed out an inappropriate applicant—but no human will have watched the weeding.

This new asymmetry—individuals with better access to legal resources (including AI itself), companies operating with increasingly self-created opacity—amplifies the dynamics that spark and

fuel lawsuits. That trend suggests three things to watch:

What happens when AI makes hiring less transparent? If a candidate feels like they're screaming into a void, they will eventually find a megaphone. Chad used a Data Access Request under the New Hampshire Privacy Act. This costs nothing. Expect more of this.

Referring to the novelty of AI sounds like a lame excuse, not a defense. We love the idea of AI "leveling the playing field" (e.g., comparing a 3.8 at Harvard to a 3.8 at a state school). But if the algorithm is "opaque" (a direct quote from an AI auditor in the report), you aren't leveling the field; you're just moving the field to a different continent, and nobody knows the rules. Aspirational language ("But we were trying to do a good thing") is not usually going to be a defense to acts that otherwise violate the law.

This is important: don't think that what AI does is literally unknowable. It isn't: Chad Markey proves it. If Chad can know it, you can know it, and enforcement agencies and juries will expect you to know it, too.

Will the burden of proof shift de facto to employers? It's fundamental in a lawsuit: the plaintiff has the burden of proving their claim. When an employment discrimination plaintiff alleges that an employer terminated her because of her sex, the employer offers—doesn't have to prove—a legitimate, non-discriminatory reason (say, poor performance or a business downturn). Then it flips back to the plaintiff to prove—yes, prove—that the reason the defendant offered is a cover for the real reason, a pretext. But it's not hard to see the plaintiff's bar shaping cases around this: if an employer can't even explain how a decision was made, then why should a jury believe an applicant was passed over for good reasons rather than bad ones?

Class actions. Always save the best for last. A class action isn't about thousands of individual grievances; it's about a single, uniform choice made by a defendant. And the "class" allegation is specifically that this choice—whether to hide gotchas in credit card agreement fine print or to pollute water—both violates the law and hurts every injured person in the same way. Individual differences in class members—their jobs or where they live, for example—don't matter; it's the legal injury that unites them.

I almost don't have to write this paragraph: now imagine the class framework applied to an AI-mediated application process that categorically excludes, say, everyone of a certain ethnic subset because—if you can't follow this leap, that's the point—the algorithm saw a relationship between shorter job tenures and a previously-unnoticed correlation to support for candidates who always vote "yes" in favor of extending federal nutrition subsidies.

What Should Employers Do Now?

If you're using AI to "filter," "normalize," or "badge" your applicants, you need to know what's happening under the hood. If a medical student can reverse-engineer your hiring process in his dorm room, imagine what a class-action firm with a budget can do. If it's not clear already, the "user clicked too fast" defense is embarrassing. And it would make every juror want to boycott you, not buy your products and services.

Stay human out there, and please contact a member of the [Kelley Drye Labor and Employment Team](#) if you have questions or worries.