

Ethikos Volume 38, Number 2. April 01, 2024

AI and automated hiring systems in AI: Emerging ethical and compliance issues

by John Rood

John Rood (john@proceptual.com, [linkedin.com/in/johnrood1/](https://www.linkedin.com/in/johnrood1/)) is the founder of Proceptual in Chicago, Illinois, USA.

It has been amazing to see how artificial intelligence (AI) has, in roughly a year, become such an engaging and important issue in our society. What modern AI can produce often feels like magic. That said, as compliance professionals, we will be on the front lines of ensuring the AI revolution is managed safely and fairly. One of the key areas where AI will be regulated is human resources (HR) and hiring. As a starting point in the discussion, here are what I consider the three core pillars of safe AI development and deployment; you'll likely hear a lot about these issues from thought leaders and government regulators.

First pillar: Transparency

Simply put, transparency refers to helping users of an AI system understand that it is in operation, what it is doing, and what data it collects. This makes sense; as a starting point to for fair use of AI, those whose job applications will be judged by the system must know what system is being used, how, and why.

We have yet to review meaningful current or draft AI regulations that do not have specific transparency or disclosure requirements. Specific to the HR context, regulations will likely require a hiring company to disclose to applicants that an AI system is in action. We generally also see requirements that job applicants be able to opt out of using the AI system, though in practice, we have not seen applicants opt out with any frequency.

Second pillar: Bias mitigation

If you ask an HR or compliance professional the first word that comes to mind when they think about AI, the answer is frequently "bias." Bias in hiring is *always* a significant issue, but what does it mean in the AI context?

In the AI hiring context, bias refers to an algorithm using protected classifications to make hiring decisions. The prototypical example of AI bias in hiring happened at Amazon nearly 10 years ago.^[1] They developed an algorithm to evaluate applicants for software engineering roles. At that time, those roles were overwhelmingly filled by men—so the algorithm started favoring men. Of course, the humans in charge told the algorithm to stop considering gender," so the algorithm started choosing factors like "played football in high school" or "did not attend an all-women's college"!

As the example shows, in no case with which I'm familiar was AI bias explicitly and purposefully programmed into a system. Instead, when AI systems are trained, they are generally trained with existing hiring data. When that data exhibits bias—as we have learned historically it often does—the algorithm itself can potentially replicate that bias.

This document is only available to subscribers. Please log in or purchase access.

Purchase Login