



Algorithmic Bias in Hiring: Fact or Myth?

 Michael A. Campion

 07-03-2025

The general public often questions the use of artificial intelligence (AI) to assist in making hiring decisions in organizations. Examples include screening applications and resumes, finding and encouraging candidates to apply, or even scoring answers to automated interviews. Considerable research, much of it conducted at Purdue's Mitch Daniels School of Business, has now been amassed to address this issue based on scientific knowledge.

Questions to ponder:

1. What is the basic issue? Why are people concerned?

The suspicion is that AI will only hire candidates who share the same demographic features as current employees or somehow limit opportunities for all candidates to compete fairly regardless of backgrounds. The concern has been labeled “algorithmic bias.”

2. What is the evidence to date?

A recent review of the 40 most-cited articles out of the 100 or so on the topic revealed that the literature has been mostly written by observers who are speculating with little knowledge of the technology. There is very little evidence that it is a widespread problem, but it could be.

3. Why might it occur?

There are some real possibilities. For example, AI algorithms are often “trained” on past candidates and employees. If the most successful candidates in the past have had backgrounds such as specific degrees, schools, past employers, work experiences, or other typical information collected on candidates, the model will select candidates with those features. This would mean those without those backgrounds would be less likely to be selected. Note that algorithms never score gender, race or other demographics directly.

4. Why might it be a myth?

One obvious possibility is that differences in hiring outcomes may be due to true differences in skills, not bias. Some demographic groups may simply have more of the job-related skills. For example, in the past men typically had more math skills and women more verbal skills.

5. What does the research indicate so far?

There have now been about 20 scientifically sound studies in personnel selection. They discovered that algorithms trained on past candidates will show the same level of subgroup differences as the data upon which they were trained, but they will not increase subgroup differences. Efforts to mathematically reduce the subgroup differences also reduce the value of the hiring procedure in predicting future job performance (called validity).

However, some very provocative new research shows that AI may be able to measure additional skills where subgroup differences are smaller, thus reducing subgroup differences and also increasing validity at the same time. This is possible because AI can analyze textual or verbal data (using Natural Language Processing), which has not been as possible in the past.

6. What should a responsible organization do?

Some methods touted as important do not matter, but might be good due diligence anyway. For example, simply ensuring that training samples have representation of all diversity subgroups may not matter, but is recommended. A more important recommendation is to examine the algorithm carefully to determine exactly what it scores and ensure it is job-related. Proof of job-relatedness is the primary legal defense to allegations of discrimination. Although some AI is thought to be “black box,” meaning its inner workings cannot be fully understood, that is probably an overstatement because it usually can be to a large extent.

Another important recommendation is to monitor the hiring outcomes of using the AI. When diversity differences are present, ensure that they are due to job-related reasons.

Yet another recommendation is to recruit better candidates from all diversity backgrounds. This is not the same as simply recruiting more diversity candidates, but instead recruiting diversity candidates who likely have the same level of skill (e.g., same schools, same

degrees). Finally, communicate what the AI measures to candidates, hiring managers and other stakeholders, because understanding is the best way to reduce suspicion.

7. What is the bottom line?

Do not assume algorithmic bias is a fact, but do not assume it is a myth either. It could be true or not in a given context. If subgroup differences in hiring outcomes are occurring, ensure they are job related. The use of AI to assist in making hiring decisions is a permanent change in modern management because of the substantial efficiencies, but we are still learning how to do it fairly and accurately. Someday the question will not be whether to use AI, but why aren't you using AI?

[Michael Campion < https://business.purdue.edu/directory/bio.php?username=campionm](https://business.purdue.edu/directory/bio.php?username=campionm)

username=campionm is the Herman C. Krannert Distinguished Professor of Business in the Department of Organizational Behavior and Human Resource Management at the Daniels School. His research spans such topics as employment testing, interviewing, mitigating employment discrimination, job analysis, work and team design, training, turnover, promotion, compensation and artificial intelligence for employment decision making.

Thought Leadership from Purdue's Business School

DANIELS INSIGHTS

[VIEW MORE INSIGHTS](#)

[SUBSCRIBE](#)

Opinions expressed on this site do not represent the official views of Purdue University or the Mitch Daniels School of Business.