

AI Isn't the Problem, We Are

The Human Roots of Bias in Artificial Intelligence

Hetansh Haria

Student

Anand Niketan Satellite School

Ahmedabad, India

Abstract—Artificial Intelligence (AI) is becoming an important part of daily life, from hiring employees to facial recognition software. Although AI is often viewed as objective, these systems frequently reflect the same biases present in the data they are trained on. This paper argues that when AI makes biased decisions, it is not the fault of technology alone, but the society that creates and provides the data used. Through examples including Amazon's hiring system and facial recognition software, this paper explores how human prejudice and inequality become embedded within artificial intelligence systems.

Index Terms—Artificial Intelligence (AI), AI bias, Algorithmic bias, Gender discrimination in hiring, Recruitment algorithms, Facial recognition bias, Racial bias in AI, Ethical AI, AI accountability, Human-generated data, Data bias, Social inequality, Automated decision-making, Fairness in AI.

I. INTRODUCTION

Can machines really be fair, or do they carry the same biases as humans? Artificial Intelligence (AI) is becoming an important part of our daily lives, from hiring employees to facial recognition software. While technology is often seen as neutral or objective, AI systems often reflect the same biases present in the data they are trained on. This makes AI bias a serious issue, as it can lead to unfair decisions that affect people's lives. In this essay I will argue that when AI makes biased decisions, it is not the fault of technology alone, but the society that creates and provides the data used. Therefore, society is ultimately responsible for the biases present in artificial intelligence systems.

II. AMAZON'S AI HIRING SYSTEM

Imagine you work on your CV for weeks only to get rejected by a computer because of your gender. In 2015, Amazon built an AI-powered system which could categorize job applications. It was trained to identify the top five applicants based on merit and qualifications to join the Amazon team. The model was trained by observing patterns in resumes submitted to the company over a 10-year period. But because most of the applicants were men, the system taught itself that male candidates were preferable, which led to female applicants getting rejected by default [1].

In this scenario, it is important to think about who is responsible. When AI makes a biased decision, society is ultimately at fault as people are the ones who create biases in the first place. If there was not a division between girls and boys, then this would not have happened. Divisions in society have always existed. People have made borders to differentiate between communities, and those communities fight over resources.

The Amazon case shows a clear gender bias. The employees previously hired at Amazon became the data used to train the AI system. This suggests that, for many years, Amazon has favoured male applicants. AI is not to blame because it is simply fed data; it is the responsibility of people working in these companies to ensure that their hiring practices are fair, and that everyone is considered for the job regardless of gender. Thus, Amazon is responsible, but not the only one at fault. Biases are not created by companies but by society as a whole.

III. FACIAL RECOGNITION AND RACIAL BIAS

Taking another example, facial recognition software made by IBM, Microsoft, and China's Megvii was tested by Joy Buolamwini of the MIT Media Lab and Timnit Gebru of Microsoft Research. They found that all three companies were more likely to correctly identify a subject's gender if that subject had pale skin [2].

To make this test fair and meaningful, they compiled and tested a dataset of 1,270 facial images sourced from lawmakers in countries where women hold a significant share of political power. Dark-skinned women were frequently misidentified as men. This happened because facial recognition systems are trained on large image datasets, and when those datasets are mostly composed of white faces, the algorithms become far more accurate at recognizing lighter skin tones while performing poorly on darker ones. The root cause is the same as in the Amazon case. The data reflected a world where darker-skinned people were underrepresented, and the AI learned accordingly.

In a facial recognition test used to determine gender, the software achieved over 99% accuracy when identifying light-skinned men. However, for darker-skinned women, its accuracy dropped significantly, with errors occurring in nearly one out of every three cases. The more confidently an AI identifies someone, the more confidently it can make decisions about them, like what ads they see or whether they are denied a service. Those decisions can still be shaped by the same societal prejudices that created the problem in the first place.

People who create AI systems, including engineers, companies, and researchers, do have serious responsibilities because their technology can affect millions of lives. They are the ones facilitating the spread of potentially incorrect information. Companies which develop AI tools must take responsibility if there is an error by the AI system. But the ultimate responsibility traces back to the humans producing the data.

IV. LARGE LANGUAGE MODELS AND HUMAN DATA

The problem of bias also runs deeper than individual systems like Amazon's hiring tool. Large language models, which are the technology behind tools like ChatGPT, work by learning patterns from vast amounts of human-written text. When you ask a question, the model does not actually know the answer like a human. They predict the most likely response based on what they have been trained on. This means that if the data contains bias, the AI will repeat it confidently.

Some of this data comes from open-source sites like Wikipedia, where anyone can contribute, meaning the information may or may not be reliable. The BBC found that around 45% of AI news queries to tools like ChatGPT, Copilot, Gemini, and Perplexity produced errors [3]. When a large language model gets something wrong, it is often because of the sources the information came from. The problem traces back to humans, and the flawed or biased content they put into the world.

V. CONCLUSION

Ultimately, when AI makes a mistake, the blame must be traced back to the society that shaped it. As seen in Amazon's hiring system and the facial recognition software developed by IBM, Microsoft, and Megvii, AI absorbs the biases that already exist in human-generated data. The data these systems learned from reflected a world where women were overlooked in hiring and darker-skinned faces were underrepresented in photographs, and these were inequalities that existed long before any algorithm was written.

Developers and companies must audit their data, question their assumptions, and be held accountable when their systems cause harm. However, fixing the technology alone is not enough. As long as the world that generates the data remains unequal, the AI trained on it will remain unequal too.

AI holds enormous promise, but it is only as fair as the society that builds it. If we want machines to treat people equally, we must first do that ourselves.

VI. ACKNOWLEDGMENT

The author would like to thank the researchers and journalists whose work on AI ethics and algorithmic bias contributed valuable insight into this topic.

REFERENCES

- [1] J. Dastin, "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women," Reuters, Oct. 2018.
- [2] J. Buolamwini, "If You're a Darker-Skinned Woman, This Is How Often Facial-Recognition Software Decides You're a Man," MIT Media Lab, Feb. 2018.
- [3] BBC News, "Largest Study of Its Kind Shows AI Assistants Misrepresent News Content 45% of the Time – Regardless of Language or Territory," BBC, Oct. 2025.