Case Study 2

Biased Amazon Recruitment Algorithm – A Lesson on AI Bias

Abstract (Summary):In 2014, a team of machine learning specialists at Amazon began working on an innovative AI-powered recruitment tool. The goal of the project was to automate the resume screening process and identify top candidates for technical roles, including software engineering positions. The system rated applications on a scale from one to five stars—similar to how customers rate products on Amazon's platform. By 2015, the team discovered a serious issue: the algorithm was systematically biased against women applying for technical roles. The tool had been trained on résumés submitted to the company over the previous decade—most of which came from men, reflecting the male dominance in the tech industry. As a result, the algorithm "learned" that male candidates were preferable. The system penalized résumés that included the word "women's" (e.g., "women's chess club captain") and downgraded the scores of graduates from two all-women's colleges. It also favored applicants who used "masculine" verbs such as "executed" and "captured" in their résumés. Despite Amazon engineers' attempts to fix the bias, the company ultimately abandoned the project in 2018, acknowledging there was no guarantee the system wouldn't develop other forms of discrimination. This case became a crucial lesson for the tech industry on the risks associated with algorithmic bias and the importance of responsible AI development.

Theme of the Case Study:

- Artificial Intelligence in Recruitment
- Algorithmic Bias
- AI Ethics
- Gender Equality in Technology

Relevance to the Reader: Amazon's recruitment algorithm serves as a key case study in algorithmic bias in AI systems. It demonstrates how machine learning systems can not only reflect but also amplify existing societal inequalities when trained on biased historical data. At a time when AI is increasingly used in hiring processes (55% of HR professionals in the U.S. predict that AI will become a routine part of their work within five years), understanding the potential risks of algorithmic bias is vital for both employers and job seekers. The Amazon case shows that even the most technologically advanced organizations can unintentionally build discriminatory systems. The case study highlights the importance of diversity and inclusion in AI development. It illustrates how a lack of diversity in tech teams and training data can lead to systems that favor one demographic group at the expense of others. For professionals working in AI ethics, this case provides a concrete example of the challenges involved in ensuring algorithmic fairness. It underscores the need for rigorous testing of AI systems for potential biases before deployment, especially in high-stakes contexts like employment. This story is also relevant in the context of emerging AI-related regulations, such as New York City's anti-discrimination law on recruitment algorithms, which came into force in 2023. It emphasizes the need to establish standards and regulations to ensure that automated decision-making systems are fair, transparent, and accountable.