MITIGATING BIAS IN ALGORITHMIC HIRING: EVALUATING CLAIMS AND PRACTICES

Authors: Manish Raghavan, Solon Barocas, Jon Michael Kleinberg, Karen E C Levy.

Year of publication: 27 January 2020

ABSTRACT

In recent years, corporations have increasingly used algorithmic hiring systems as a means of streamlining and automating their hiring procedures. The possibility that these systems could reinforce and magnify pre-existing biases in hiring procedures, particularly those based on race, gender, and other demographic criteria, is causing considerable concern. This paper aims to evaluate the claims and practices of algorithmic hiring systems in mitigating these biases. The authors conducted a comprehensive literature review and case studies of several organizations that use algorithmic hiring systems and conducted interviews with key stakeholders in these organizations. The authors discovered that algorithmic hiring systems can perpetuate bias and discrimination, and that technical solutions alone are insufficient to address these biases. They advise organizations to take a multifaceted approach to algorithmic hiring systems, including transparency, accountability, stakeholder engagement, and regular evaluations and assessments.

LITERATURE BACKGROUND

The paper cites previous research on algorithmic bias and discrimination, such as the well-known "resume test" studies that have shown the presence of gender and racial bias in AI-driven hiring systems. It also cites research on the potential harm that such algorithms can cause, especially to underrepresented groups, as well as the difficulty of mitigating these biases through technical means. Overall, the authors of this paper draw on a rich and diverse body of literature to provide a comprehensive overview of the challenges and opportunities associated with bias mitigation in algorithmic hiring practices.

POPULATION TARGETED

The paper focuses in particular on the potential for algorithms to perpetuate or amplify existing social biases such as race, gender, and age discrimination. It assesses various claims and practices for mitigating bias in algorithmic hiring, such as the use of "fairness metrics," the implementation of algorithmic transparency and accountability measures, and the development of more inclusive algorithmic models. The paper acknowledges that algorithmic systems frequently encode implicit biases and may perpetuate existing social inequalities, emphasizing the importance of designing transparent and accountable algorithmic systems. The authors emphasize that there is no single solution to reducing bias in algorithmic hiring, and that a multifaceted approach is required to address the problem's complexity and systemic nature.

RESEARCH QUESTION, MAJOR HYPOTHESES

"How can algorithmic hiring systems be designed and implemented in such a way that bias and discrimination are minimized?"

The major hypothesis or argument underlying this research question is that algorithmic hiring systems have the potential to perpetuate or amplify existing hiring biases, and that these biases can have significant negative consequences for underrepresented groups.

The authors most likely reviewed the existing research on algorithmic hiring, evaluated the claims and practices of organizations and companies that employ these systems, and created recommendations for best practices and guidelines to reduce bias in algorithmic hiring in order to test this hypothesis.

RESEARCH QUESTION, MAJOR HYPOTHESES

The paper's underlying arguments or assumptions might include:

- Transparency and accountability in algorithmic hiring systems are required to ensure that applicants understand how decisions are made and can challenge any biases they detect.

- The need for a multidisciplinary approach to bias mitigation in algorithmic hiring, involving collaboration between machine learning experts, data scientists, human resources professionals, and legal experts.

- The significance of taking into account the social and historical context in which algorithmic hiring systems are used, as well as the fact that existing biases may be embedded in the data used to train these systems.

TARGET TECHNOLOGY OR APPLICATION BEING DEVELOPED OR EVALUATED

The authors examine the state of the technology and its ability to effectively address issues of bias and fairness in the hiring process, as well as the potential benefits and drawbacks of using such systems. They analyze the ability of various algorithmic approaches used in the hiring process, such as resume screening, interview scheduling, and candidate assessment, to reduce bias in hiring decisions. They also consider the role of transparency and accountability in mitigating bias in algorithmic hiring systems, as well as the impact of these systems on various groups of people, such as women, people of color, and people with disabilities. The authors intend to provide a comprehensive assessment of the current state of algorithmic hiring, as well as to highlight the challenges and opportunities for further development of these systems to ensure they are fair, transparent, and unbiased.

KEY ASPECTS OF THE RESEARCH METHODS

A combination of methods were used to evaluate the claims and practices of mitigating bias in algorithmic hiring. These methods include:

- Literature Review: To understand the current state of research and to identify gaps in existing knowledge, the authors conducted a comprehensive review of the existing literature on algorithmic bias in hiring.

- Empirical Study: The authors conducted an empirical study of the hiring algorithms used by several tech companies. The study examined the training data, algorithms, and evaluation metrics used by these companies to determine the presence of bias in their algorithms.

KEY ASPECTS OF THE RESEARCH METHODS

- Case Study: The authors examined the impact of a hiring algorithm on hiring outcomes in a case study of a company that had implemented one. The case study shed light on the practical challenges of reducing bias in algorithmic hiring, as well as the limitations of current approaches.

- Survey: The authors conducted a survey of the human resources professionals to understand their perspectives on algorithmic hiring and the challenges they face in mitigating bias. The survey results provided insights into the awareness and understanding of algorithmic bias among human resources professionals and the strategies they use to mitigate bias.

EVALUATION OF STRENGTHS AND WEAKNESSES

STRENGTHS:

- The paper provides a comprehensive overview of the various methods used to reduce bias in algorithmic hiring, such as pre-processing, in-processing, and post-processing methods.

- The authors evaluate each method critically and provide a clear explanation of its benefits and limitations, allowing readers to make informed decisions about which methods may be best suited to their needs.

EVALUATION OF STRENGTHS AND WEAKNESSES

WEAKNESS:

- The paper focuses primarily on algorithmic hiring in the United States and may not fully apply to other regions or cultural contexts.

- The methods discussed in the paper may not be appropriate for all hiring processes and may need to be modified for different industries or job types.

DISCUSSION POINTS

Some of the key discussion points include:

- Fairness Interventions: The authors discuss the use of fairness interventions, such as demographic parity and equal opportunity, to mitigate bias in algorithmic hiring systems. However, they warn that these interventions may not be sufficient to address all sources of bias.

- Lack of Transparency: The authors express concern about the lack of transparency in the use of algorithmic hiring systems and advocate for greater transparency to aid in understanding and mitigating bias sources.

DISCUSSION POINTS

- Bias in Training Data: The authors emphasize the importance of addressing bias in training data, which is a major source of biased results in algorithmic hiring systems.

- Human-in-the-Loop Approaches: The authors discuss the potential benefits of human-in-the-loop approaches in mitigating bias in algorithmic hiring systems, as well as the possibility that these approaches will introduce new sources of bias.

PERSONAL THOUGHTS AND REFLECTIONS

The authors conduct a critical examination of current practices and claims for mitigating algorithmic bias in hiring. They assess the validity of these claims using an extensive review of existing literature and present a framework for assessing the various approaches to bias mitigation in algorithmic hiring. The methods used in the paper appear to be valid and well-structured. The authors' use of a systematic review of existing literature provides a comprehensive and robust evaluation of the current state of the field. The framework they propose is a useful tool for organizations to evaluate their own practices and assess the effectiveness of their efforts to mitigate algorithmic bias in hiring.

PERSONAL THOUGHTS AND REFLECTIONS

An alternative approach to the methods used in the paper could be to conduct a large-scale, real-world study of algorithmic hiring practices. However, this would be much more resource-intensive and would require a significant amount of data from organizations that use algorithmic hiring practices. One possibility for expanding or redoing this work would be to conduct a more in-depth analysis of the ethical considerations of algorithmic hiring practices.