AI for Algorithmic Auditing: Mitigating Bias and Improving Fairness in Big Data Systems

Priya Patel
Department: Department of Computer Science
University: Indian Institute of Technology Delhi (IITD), India
priya.patel@iitd.ac.in

Mohd Nasim Uddin
Teesside University
nasimuddin2011@gmail.com

Abstract

Algorithmic decision-making systems are being increasingly used in high-impact domains like finance, healthcare, and criminal justice. However, these systems can unintentionally discriminate against certain groups due to biases in training data or models. This has led to calls for increased transparency and algorithmic auditing to detect and mitigate unfairness. This paper provides an overview of emerging techniques using AI to audit black-box systems for bias. First, we discuss sources of algorithmic bias and the importance of fairness in AI systems. We then review different definitions and metrics for fairness, including group versus individual notions of fairness. Next, we survey different algorithmic auditing methods to assess system behavior using only input/output queries. These include techniques based on causality, counterfactual reasoning, and adversarial models. We also examine methods to improve system fairness by detecting and mitigating biases in the training pipeline, modifying model parameters directly, or post-processing model outputs. Finally, we outline key challenges and opportunities, including model interpretability, scalability, and the need to incorporate domain expertise into notions of fairness. Overall, this paper synthesizes recent advancements in using AI for algorithmic auditing and provides insights into translating these methods into practice to build more trustworthy and ethical AI systems.

Indexing terms: Algorithmic decision-making, AI systems, Bias, Fairness, Auditing methods

Introduction

The expansion of algorithmic decision-making systems, driven by artificial intelligence (AI) and machine learning technologies, has become increasingly pervasive across critical sectors like finance, healthcare, criminal justice, and human resources. These systems, designed for predictive analytics and automated decision-making, have raised significant concerns regarding the presence of inherent biases that may lead to unintended discrimination. The inadvertent biases can be attributed to various factors, including historical imbalances, non-representative training data, and limited model generalizability. Instances of bias have been observed in different contexts, such as resume screening algorithms displaying preferences against candidates from women's colleges and healthcare algorithms exhibiting racial bias in estimating clinical risk scores [1]. The repercussions of algorithmic bias are profound, extending to the tangible impact on individuals' lives. Biased decision-making can perpetuate historical discrimination, reinforcing existing disparities in opportunities and outcomes. In addition to its societal implications, algorithmic bias can also contribute to a erosion of trust in AI applications. When individuals perceive that automated systems are making decisions based on unfair or discriminatory criteria, confidence in the technology diminishes, hindering its widespread acceptance and adoption [2].

Figure 1.

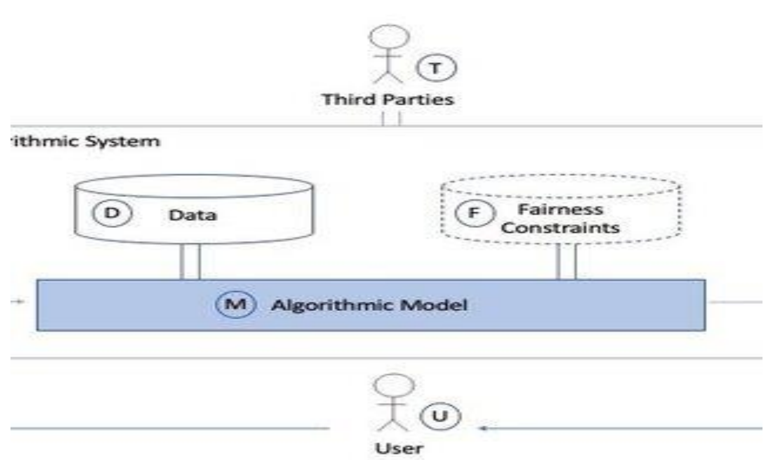

Addressing algorithmic bias requires a multi-faceted approach. Firstly, it involves meticulous scrutiny of training data to identify and rectify biases that may have inadvertently been ingrained during the learning process. Ensuring that training datasets are representative of diverse demographics is crucial in mitigating bias [3]. Moreover, continuous monitoring of AI systems in real-world applications is essential to detect and rectify biases that may emerge over time as the system encounters new data. Transparency in the decision-making process is another critical element. Providing stakeholders with insights into how algorithms arrive at their decisions fosters accountability and allows for external validation.

Addressing these challenges necessitates a shift in evaluating algorithmic systems, especially when source code and models are proprietary, making traditional software testing methods impractical. The prevailing approach involves treating systems as black-box entities and relying on input/output queries for assessment. This method, described as practical for real-world systems, involves probing with meticulously crafted test inputs to analyze outputs, thereby detecting biases, measuring fairness, and diagnosing issues solely based on observable behaviors. The demand for transparency and accountability has spurred the development of algorithmic auditing techniques that leverage AI and machine learning [4]. These techniques, incorporating insights from causality, counterfactual reasoning, and adversarial methods, aim to facilitate rigorous, scalable, and continuous auditing of black-box systems. The overarching objective is to establish Trustworthy AI systems characterized by ethicality, fairness, and audibility. Algorithmic auditing not only offers a means to scrutinize system behavior but also provides valuable feedback to refine the training process and enhance the models, particularly in high-stakes domains where understanding model decisions, fairness, and risks is as pivotal as predictive accuracy.

This paper comprehensively reviews algorithmic auditing techniques utilizing AI to assess and enhance the fairness of black-box systems. Section 2 delves into the background, elucidating sources of bias and emphasizing the imperative of mitigating discrimination in AI systems. Section 3 subsequently expounds upon diverse formal definitions and metrics for fairness [5]. Moving forward, Section 4 surveys emerging techniques for auditing black-box systems, exclusively relying on input/output analysis. Section 5 explores algorithms designed to mitigate bias and enhance fairness post-detection of issues. Finally, Section 6 concludes by delving into the key challenges and opportunities associated with translating algorithmic auditing methods into practical applications.

Background on Algorithmic Bias

Algorithmic bias, a phenomenon gaining heightened scrutiny in the realm of artificial intelligence, pertains to instances where automated decision-making systems yield unintentionally prejudiced or unfair outcomes, resulting in discrimination against specific groups. The manifestation of such bias occurs without malicious intent and may surface even when protected attributes like gender or race are not overtly factored into the algorithms. The escalating real-world influence of AI systems, coupled with the opacity surrounding their decision-making processes, has intensified concerns regarding bias. Despite the prevailing expectation for algorithms to exhibit neutrality and objectivity, several factors within the modeling pipeline can surreptitiously introduce bias, complicating the pursuit of fairness and impartiality [6].

The intricacies of algorithmic bias necessitate a thorough examination of the various stages within the modeling pipeline where biases may emerge. One critical point of consideration is the training data used to develop and fine-tune algorithms. If the training data is inherently biased, reflecting historical disparities or societal prejudices, the algorithm may inadvertently perpetuate and amplify those biases in its decision-making. Addressing this challenge requires meticulous scrutiny of training datasets and the implementation of corrective measures to mitigate existing biases or prevent their reinforcement [7]. Moreover, the algorithms themselves, designed by human developers, may inadvertently encode biases based on the creators' perspectives, beliefs, or assumptions. Unintentional biases can be ingrained in the choice of features, the formulation of decision rules, or the optimization processes. To counteract this, a rigorous evaluation of the algorithmic design, including a comprehensive audit for potential biases, becomes imperative. Implementing measures such as diverse and inclusive development teams, alongside continuous monitoring and auditing, can contribute to reducing biases at the algorithmic level [8].

Another critical aspect is the interpretability of AI models, as the lack of transparency in how algorithms reach specific decisions can exacerbate concerns about bias. Black-box algorithms, which operate without providing clear explanations for their outputs, hinder the identification and rectification of biased patterns. Enhancing the interpretability of AI models through techniques like explainable AI (XAI) can facilitate a better understanding of the decision-making process, allowing for the detection and correction of biased outcomes. The evolving nature of societal norms and ethical standards further complicates the challenge of addressing algorithmic bias. What may be considered biased today might not be perceived similarly in the future. Continuous engagement with diverse stakeholders, including ethicists, policymakers, and affected communities, is essential to establish dynamic frameworks for evaluating and mitigating bias in AI systems. Ethical considerations must be an integral part of the development lifecycle, and mechanisms for ongoing reflection and adaptation should be institutionalized.

First, training datasets themselves can be biased, reflecting historical discrimination or failing to adequately represent certain groups. Labels can also be biased, subjective, or incorrectly measured for underrepresented groups. Models trained on such data inherit and propagate these biases. Second, the choice of model, assumptions, and features can lead to biased behavior in minority groups not sufficiently covered in training. Finally, biases can be introduced when deploying models to new environments and populations different from training. Even without explicit protected attributes, models can still exploit proxy variables correlated with race or gender, for example. Left unchecked, algorithmic bias can have serious detrimental impacts on people's lives - denying opportunities, resources, and information. It also further marginalizes vulnerable groups already facing structural disadvantages [9]. There are compelling ethical arguments around principles of justice, fairness, and preventing discrimination. Practically, biased systems also undermine public trust in AI which can dampen adoption and innovation. And anti-discrimination laws prohibit algorithmic discrimination in many contexts like hiring, lending, and public services.

These concerns have led to increased focus on algorithmic accountability and transparency to understand sources of bias and ensure fair outcomes. However, most real-world systems rely on proprietary data and complex models like deep neural networks. Classical software testing with full code and model access is infeasible. Algorithmic auditing has emerged as a crucial methodology to still analyze real systems deployed "in the wild" using only input/output queries. This black-box perspective audits systems based solely on their observable behaviors rather than internal details. Next, we survey different definitions and metrics to assess fairness in algorithmic systems.

Definitions and Metrics for Fairness

Fairness is an intricate, multi-dimensional concept that lacks consensus definitions and metrics. This poses challenges in formulating auditing approaches and improving systems. Broadly, notions of fairness judge whether a system's outcomes are equitable across different groups based on sensitive attributes like gender or race. However, systems can be fair according to one criterion but unfair to another. Important distinctions also exist between group notions of fairness versus individual fairness. Here we overview common definitions and measures that underlie approaches for auditing and mitigating unfairness.

Many definitions consider group fairness - the treatment of different protected groups defined by sensitive attributes. Demographic parity requires the overall proportion of positive outcomes be equal between groups. For example, lending approvals should be equal across ethnicities [10]. However, this can conflate different base rates of risk across groups. Other notions like equalized odds and equal opportunity focus specifically on true/false positive/negative rates across groups. Equalized odds require equal true positive rates (TPR) and false positive rates (FPR) between groups. Equal opportunity equalizes the TPR but allows FPR differences. Criteria based on TPR/FPR discrepancies are common in contexts like hiring or lending where erroneous outcomes have asymmetric costs.

Fairness definitions also consider individual notions centered on similar individuals receiving similar outcomes. Individual fairness requires a similarity metric between individuals and that similar individuals under this metric receive similar outcomes.

While attractive conceptually, individual fairness relies heavily on appropriately defining individual similarity for the application domain. It also does not consider representation discrepancies between groups. There are also statistical frameworks like causality to formalize discrimination. Counterfactual fairness requires model outcomes be invariant to changes in protected attributes. Intersectional fairness considers compound biases against individuals with multiple sensitive attributes. Game theoretic definitions of fairness based on equilibrium concepts have also emerged recently.

This diversity of competing notions poses challenges in operationalizing fairness. Systems satisfying one definition may still be biased per another. Picking appropriate, context-specific definitions requires understanding trade-offs between metrics and incorporating expertise of domain impacts and constraints. Furthermore, few definitions fully capture the intricate real-world dynamics of bias. Moving forwards, auditing systems against multiple definitions provides a more comprehensive perspective. Fairness should be viewed as satisfying multiple, potentially conflicting dimensions rather than achieving one singular objective [11].

With these considerations in mind, next we examine algorithmic auditing techniques to assess fairness of black-box systems using only input/output analysis.

Table 1: Comparison of different definitions and metrics for algorithmic fairness

Definition	Key Concept	Advantages	Limitations
Demographic Parity	Equal outcomes proportions between groups	Intuitive, clear meaning	Conflates underlying risk/qualification differences
Equalized Odds	Equal TPR/FPR between groups	Considers asymmetric error costs	Ignores overall population impacts
Individual Fairness	Similar individuals have similar outcomes	Judges individuals, flexible similarity metric	Hard to define appropriate similarity, no group notions
Counterfactual Fairness	Invariance to changes in protected attributes	Uses causal reasoning	Computationally intensive

Black-box Auditing Techniques: Auditing the fairness of real-world systems requires viewing them as black-boxes since details like models and data are often proprietary. We rely solely on querying systems with inputs and analyzing outputs to audit their behaviors. Here we review emerging techniques to audit black-box systems for algorithmic fairness.

Causality-based Approaches: Techniques based on causal reasoning offer robust ways to audit systems and measure bias and discrimination. Causal models explicitly capture relationships between sensitive attributes, intermediate variables, and outcomes. This allows distinguishing correlations from causation to identify legitimate and illegitimate discrimination through auditing. Kilbertus et al. develop a causal auditing framework where systems are represented as causal models relating sensitive attributes A to outcomes Y through observed proxies X. Bias is quantified by estimating effects along unfair causal pathways from A to Y that should be eliminated to ensure fairness [12].

Counterfactual Reasoning: Counterfactual reasoning considers how outcomes would change under different input conditions to discern biases. Wexler et al. formulate a framework for auditing based on estimating counterfactuals - the outcomes that would have occurred for an individual if their protected attributes differed [13]. For example, would a loan applicant have received a different decision if their gender was different? Comparing counterfactuals across individuals reveals biased treatment. Estimating counterfactuals relies on building a predictive model from system outputs to impute outcomes.

Adversarial Techniques: Adversarial approaches generate synthetic inputs to maximally expose biases in black-box systems. Perturbed inputs are optimized to produce outcomes that most differentiate groups according to measures like demographic parity. This allows stress testing systems and quantifying fairness. Theura et al. propose an adversarial sampling technique for auditing where a discriminator model is simultaneously trained to predict sensitive attributes from the system's outputs. The

generator finds inputs to maximize prediction accuracy of the discriminator, thus revealing cases with the most bias.
Bias Testing with Constrained Optimization: Auditing can also be framed as constrained optimization problems to quantify bias. Conditional demographic parity is formalized as minimizing a loss function over outcomes subject to the constraint that an auditor model cannot reliably infer protected attributes from outputs. Optimization-based formulations offer flexibility in encoding different fairness definitions into constraints and objectives. However, solving the optimizations can be prohibitively expensive for real systems [14]. Approximation methods are needed to scale auditing.

Overall, these emerging techniques enable practical black-box auditing of algorithmic systems based on querying their input/output behaviors. Causal, counterfactual, adversarial, and optimization approaches offer complementary strengths in diagnosing and measuring different aspects of system fairness. Next, we examine techniques to mitigate unfairness once auditing exposes issues [15].

Table 2: Overview of different techniques for black-box auditing of algorithmic fairness

Method	Key Idea	Strengths	Limitations
Causal Modeling	Audit along unfair causal pathways from sensitive attributes	Robust bias identification	Assumes valid causal model
Counterfactual Estimation	Compare outcomes if inputs changed	Flexible bias quantification	Difficult outcome estimation
Adversarial Input Generation	Maximize bias through optimized inputs	Stress testing system	Not guided by formal fairness criteria
Constrained Optimization	Formal auditing as constrained optimization	Precisely encode constraints and objectives	Computationally challenging

Improving Fairness

Detecting unfairness through auditing provides crucial feedback to improve system design and modeling. We now overview techniques to mitigate biases and enhance fairness:
1. Improving training data and pipelines: Many biases stem from suboptimal data collection, labeling, and preprocessing. Strategies like smart data augmentation, selection and weighting, and relabelling can reduce representation biases and label errors. Causal modeling of data generation can also help select fairer training data and variables.
2. Modifying models and parameters: Algorithms like adversarial debiasing directly constrain model objective and gradients during training to minimize prediction of protected attributes. Controlled post-processing of embeddings is another technique to remove information about sensitive attributes. Optimizing models to satisfy formal fairness constraints defined over outputs is also possible.
3. post-processing model outputs: Even without retraining models, post-processing methods like rejecting discriminatory predictions and calibration can help satisfy fairness criteria. Output thresholds and decision rules can be adjusted to equalize metrics between groups. However, this does not address root causes and can impact overall performance.
4. Providing explanations: Highlighting influential features and providing explanations for model decisions helps build trust and reveal questionable dependencies. Users can then provide feedback to systematically improve the model. Interactive approaches have promise to incorporate human notion of fairness.
5. Incorporating ethics frameworks: Co-designing systems with ethics principles and meaningful oversight helps proactively address biases rather than retrofitting fairness. Solutions should consider unique constraints and biases of the application context. Models must be conceptualized as sociotechnical systems fraught with assumptions and tradeoffs.

A combination of the above strategies is needed to holistically address the various mechanisms by which bias arise and propagate through the modeling pipeline. Importantly, improving fairness requires incorporating domain expertise into what constitutes harms, constraints, and appropriate notions of algorithmic fairness for the application context. Overall, algorithmic auditing provides the crucial feedback loop to guide interventions that move towards ethical AI systems.

Table 3: Overview of strategies to improve algorithmic fairness.

Approach	Method	Strengths	Limitations
Improve Training Process	Data augmentation, relabeling, variable selection	Addresses data root causes	Significant workflow changes
Modify Model Parameters	Adversarial training, gradient constraints	Direct encoding of fairness	Model quality tradeoffs
Post-process Outputs	Threshold adjustments, prediction dropping	Simple implementation	Superficial, limited performance impacts
Provide Explanations	Feature importance, example based	Builds trust and user feedback	Limited scope, humans also biased
Incorporate Ethics Frameworks	Co-design, meaningful oversight	Fundamentally addresses biases and assumptions	Challenging adoption at scale

Challenges and Opportunities

While growing in maturity, algorithmic auditing and bias mitigation face key challenges translating research into widespread practice:

Complex, subjective nature of fairness: Handling the intricate, context-dependent concept of fairness poses fundamental conceptual difficulties. There is need for multi-disciplinary collaboration with law, social sciences, and ethics to formulate appropriate notions. Platforms to solicit impact categories and preferences from affected communities are also important.

Scalability: Current methods analyze models individually, limiting scalability across large, evolving systems. Approaches leveraging system similarity and transfer learning hold promise for scale. Streamlined workflows and software infrastructure are also needed for continuous auditing.

Interpretability: To enable understanding and improvement, audits should provide rich explanations of failure modes beyond binary pass/fail assessments. Generating full descriptions and examples of model biases poses open research questions.

Incorporating human oversight: Purely automated techniques have limited scope. Interactive frameworks where humans analyze high-level trends and provide feedback to algorithms can contextualize auditing. But human reviewers have their own biases which require caution.

Legal compliance: Regulations like Europe’s GDPR prohibit certain uses of sensitive personal data which seemingly complicate audits requiring such data. Interpreting compliance obligations in the context of auditing remains open. Careful system design can likely enable auditing under data protection laws.

Adoption incentives: Organizations may resist transparency that exposes issues, despite ethical imperatives. But auditing can also improve products and efficiency. Clearly conveying these mutual benefits for users and providers can enable adoption.

The opportunities for impact are tremendous if these open challenges are navigated appropriately [16]. Overall, there is clear momentum at the interface of AI, fairness, accountability, and ethics. Auditing and continuous improvement should become integral parts of building robust, transparent, and socially responsible AI systems. Technological solutions alone are insufficient, but auditing frameworks provide feedback mechanisms to actualize ethical principles in practice. This review synthesizes the landscape of emerging techniques at the nexus of AI, algorithmic auditing, and fairness.

Case Studies

It is also instructive to examine real-world case studies where algorithmic auditing and debiasing have been applied to improve high-impact systems:

Judicial Risk Assessments: Actuarial risk assessment tools are used across the United States to guide bail and sentencing decisions. While promoted as more "objective" than unaided human judgement, analyses found several prominent tools to exhibit significant racial biases, falsely flagging black defendants as higher risk. This underscores the need to rigorously audit proprietary tools rather than assuming their fairness. In response, optimized auditing algorithms were developed to expose biases in COMPAS, a widely used risk assessment tool. This enabled courts and developers to begin addressing issues. The stakeholder response showcases the complex interplay between technology, public scrutiny, and policy change in debiasing [17].

Online Advertising: Significant gender biases have been demonstrated in online advertising systems, with STEM job ads disproportionately shown to men. This likely stems from historical imbalances in industries and demographics of ad clickers. Besides being unfair, such biases propagate gender gaps further. Google adopted an internal tool to algorithmically audit ads for demographic parity, leading to policy changes and dramatic reductions in exposure bias. However, other definitions like equalized odds may be more appropriate to balance competing objectives like click-through rates. This demonstrates both the feasibility and nuances of applying auditing in large-scale production systems [18].

Healthcare Analytics: Many documented cases show medical algorithms exhibiting racial biases in estimating risk scores and treatment recommendations. Issues stem from underrepresented minorities in the training data as well as complex causal relationships between race, socioeconomic factors, and health. Simply removing racial data is inadequate and reduces personalization. To address this, one study tailored an adversarial debiasing approach to improve the fairness of sepsis prediction models without compromising accuracy. The improved predictions demonstrate the value of context-specific auditing and debiasing.

These examples highlight the progress towards fairer AI systems, as well as the considerable work remaining [19]. Thorough auditing to quantify biases in real-world systems lays the foundation. But this must be coupled with iteratively enhancing data, models, and assumptions - in collaboration with domain experts and affected communities. Overall, the reviewed landscape of techniques constitutes a promising path towards equitable and trustworthy algorithmic systems.

Conclusion

Given the pervasive influence of algorithmic decision-making in critical domains, a proactive approach is imperative to address the potential amplification of biases within these systems. The widespread adoption of algorithms in finance, justice, healthcare, and other sectors highlights the urgency of ensuring that these systems do not inadvertently perpetuate or accentuate discriminatory practices. Recognizing this imperative, the focus on auditing proprietary algorithms deployed in real-world scenarios has gained traction [20]. The scrutiny of these algorithms as black-box entities is a foundational step towards mitigating bias. Algorithmic bias can emerge from various sources, including biased training data, flawed model architectures, or implicit biases embedded in the design process. Auditing algorithms involves a systematic examination of these components to identify and rectify biases that may impact decision outcomes. The black-box nature of many proprietary algorithms necessitates specialized auditing techniques, emphasizing transparency and accountability in algorithmic decision-making [21]. This process involves deciphering the intricate workings of algorithms, understanding their decision pathways, and assessing the potential biases embedded in their decision criteria.

The financial sector, for instance, heavily relies on algorithms for decision-making processes, ranging from credit scoring to investment strategies. Auditing algorithms in finance requires an in-depth analysis of the variables considered, the weight assigned to each variable, and the decision thresholds set by the algorithm. This approach ensures that financial algorithms adhere to ethical standards and do not discriminate against individuals based on factors such as race, gender, or socioeconomic background.

Similarly, in the realm of criminal justice, algorithms are employed for risk assessment, sentencing recommendations, and parole decisions [22]. The potential biases in these algorithms can have profound implications for individuals within the justice system [23]. Algorithmic auditing in this context involves scrutinizing the training data used to develop these systems, evaluating the fairness of the underlying algorithms, and assessing the impact of these algorithms on marginalized communities. This systematic examination is essential to uphold the principles of fairness and justice in algorithmic decision-making within the legal domain [24].

In healthcare, where algorithms play a crucial role in diagnostics and treatment planning, auditing becomes paramount to ensure that these systems do not contribute to healthcare disparities. Auditing healthcare algorithms involves evaluating their accuracy, sensitivity, and specificity, while also examining the potential biases in the data used for training. This rigorous assessment helps prevent situations where certain demographic groups may receive disparate healthcare outcomes due to algorithmic biases [25]. The call for auditing proprietary algorithms extends beyond individual sectors, encompassing a broader societal need for responsible and ethical AI deployment. Governments, regulatory bodies, and industry standards organizations are increasingly recognizing the importance of establishing guidelines and frameworks for algorithmic auditing. These frameworks aim to standardize the auditing process, promote transparency, and hold organizations accountable for the ethical implications of their algorithmic systems.
This paper has provided an overview of contemporary techniques for black-box algorithmic auditing, emphasizing the utilization of input/output analysis and core AI methodologies such as causality and adversarial learning. By acknowledging the challenges posed by opaque algorithms, the research community has made strides in developing tools that facilitate rigorous and scalable audits. The significance of these audits lies not only in identifying biases but also in offering insights into potential mitigation strategies. Upon uncovering biases through auditing processes, it is essential to implement robust strategies to mitigate these issues. The paper has touched upon key strategies for bias mitigation, recognizing the critical role of adversarial learning and causality in addressing algorithmic disparities [26]. However, it is crucial to acknowledge that the landscape of algorithmic auditing is dynamic, and open research challenges persist. Issues related to the complexity of algorithms, scalability of auditing processes, and the necessity of human oversight represent ongoing areas of concern that demand further exploration. Despite these challenges, the intersection of AI, accountability, and ethics continues to evolve, holding the promise of fostering fairer automated systems. The tools and frameworks reviewed in this paper represent crucial advancements in the pursuit of Trustworthy AI. As the field progresses, addressing the complexities associated with algorithmic decision-making remains paramount, and ongoing research endeavors are essential to refining auditing methodologies, ensuring their applicability at scale, and upholding the ethical standards of AI deployment. In essence, algorithmic auditing serves as a pivotal feedback loop, contributing to the realization of AI's potential to enhance social welfare while minimizing the risks associated with biased decision-making [27].

References

[1] M. H. Jarrahi, G. Newlands, M. K. Lee, C. T. Wolf, E. Kinder, and W. Sutherland, "Algorithmic management in a work context," *Big Data Soc.*, vol. 8, no. 2, p. 205395172110203, Jul. 2021.
[2] Y. Ivanova, "The role of the EU fundamental right to data protection in an algorithmic and big data world," in *Data Protection and Privacy*, Hart Publishing, 2021.
[3] M. Kamal and T. A. Bablu, "Machine Learning Models for Predicting Click-through Rates on social media: Factors and Performance Analysis," *IJAMCA*, vol. 12, no. 4, pp. 1–14, Apr. 2022.
[4] I. Iglezakis, T. Trokanas, and P. Kiortsi, "The right not to be subject to automated individual decision-making/profiling concerning big health data. Developing an algorithmic culture," *SSRN Electron. J.*, 2021.
[5] I. Iglezakis, T. Trokanas, and P. Kiortsi, "The right not to be subject to automated individual decision-making/profiling concerning big health data. Developing an

algorithmic culture," *International Journal of Data Science and Big Data Analytics*, vol. 1, no. 3, p. 6, Nov. 2021.
[6] M. Muniswamaiah, T. Agerwala, and C. Tappert, "Big data in cloud computing review and opportunities," *arXiv preprint arXiv:1912.10821*, 2019.
[7] S. Hogue, "Project maven, big data, and ubiquitous knowledge: The impossible promises and hidden politics of algorithmic security vision," in *Automating Crime Prevention, Surveillance, and Military Operations*, Cham: Springer International Publishing, 2021, pp. 203–221.
[8] D. Agarwal, R. Sheth, and N. Shekokar, "Algorithmic trading using machine learning and neural network," in *Computer Networks, Big Data and IoT*, Singapore: Springer Singapore, 2021, pp. 407–421.
[9] Anjala, "Algorithmic assessment of text based data classification in big data sets," *J. Adv. Res. Dyn. Control Syst.*, vol. 12, no. SP4, pp. 1231–1234, Mar. 2020.
[10] S. Sioutas, G. Vonitsanos, N. Zacharatos, and C. Zaroliagis, "Scalable and hierarchical distributed data structures for efficient big data management," in *Algorithmic Aspects of Cloud Computing*, Cham: Springer International Publishing, 2020, pp. 122–160.
[11] L. Yuqi, "Analysis of algorithmic infringement risk in the background of big data," *Criminal Justice Science & Governance*, vol. 1, no. 1, pp. 88–97, 2020.
[12] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Approximate query processing for big data in heterogeneous databases," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5765–5767.
[13] S. Ahmadian and S. Haddadan, "A theoretical analysis of graph evolution caused by triadic closure and algorithmic implications," in *2020 IEEE International Conference on Big Data (Big Data)*, Atlanta, GA, USA, 2020.
[14] Y. Cui, J. Song, M. Li, Q. Ren, Y. Zhang, and X. Cai, "SDN-based big data caching in ISP networks," *IEEE Trans. Big Data*, vol. 4, no. 3, pp. 356–367, Sep. 2018.
[15] A. Nassar and M. Kamal, "Ethical Dilemmas in AI-Powered Decision-Making: A Deep Dive into Big Data-Driven Ethical Considerations," *IJRAI*, vol. 11, no. 8, pp. 1–11, Aug. 2021.
[16] J. Grosman and T. Reigeluth, "Perspectives on algorithmic normativities: engineers, objects, activities," *Big Data Soc.*, vol. 6, no. 2, p. 205395171985874, Jul. 2019.
[17] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Federated query processing for big data in data science," in *2019 IEEE International Conference on Big Data (Big Data)*, 2019, pp. 6145–6147.
[18] A. Nassar and M. Kamal, "Machine Learning and Big Data Analytics for Cybersecurity Threat Detection: A Holistic Review of Techniques and Case Studies," *Intelligence and Machine Learning …*, 2021.
[19] P. B. de Laat, "Algorithmic decision-making based on machine learning from big data: Can transparency restore accountability?," *Philos. Technol.*, vol. 31, no. 4, pp. 525–541, 2018.
[20] J. Levy and R. Prizzia, "From data modeling to algorithmic modeling in the big data era: Water resources security in the Asia-pacific region under conditions of climate change," in *Advanced Sciences and Technologies for Security Applications*, Cham: Springer International Publishing, 2018, pp. 197–220.
[21] M. Krenn *et al.*, "Creating a large-scale silver corpus from multiple algorithmic segmentations," in *Medical Computer Vision: Algorithms for Big Data*, Cham: Springer International Publishing, 2016, pp. 103–115.
[22] C. Khatri *et al.*, "Algorithmic content generation for products," in *2015 IEEE International Conference on Big Data (Big Data)*, Santa Clara, CA, USA, 2015.
[23] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Context-aware query performance optimization for big data analytics in healthcare," in *2019 IEEE High Performance Extreme Computing Conference (HPEC-2019)*, 2019, pp. 1–7.
[24] M. Hong, M. Razaviyayn, Z.-Q. Luo, and J.-S. Pang, "A unified algorithmic framework for block-structured optimization involving big data," *arXiv [math.OC]*, 09-Nov-2015.
[25] E. H. Nathanael, B. Hendradjaya, and W. Danar Sunindyo, "Study of algorithmic method and model for effort estimation in big data software development case study: Geodatabase," in *2015 International Conference on Electrical Engineering and Informatics (ICEEI)*, Denpasar, Bali, Indonesia, 2015.

[26] R. Nyman and P. Ormerod, “Big Data, Socio-psychological theory, algorithmic text analysis and predicting the Michigan Consumer Sentiment Index,” *arXiv [q-fin.ST]*, 22-May-2014.
[27] M. Winlaw, M. B. Hynes, A. Caterini, and H. De Sterck, “Algorithmic acceleration of parallel ALS for collaborative filtering: Speeding up distributed big data recommendation in spark,” in *2015 IEEE 21st International Conference on Parallel and Distributed Systems (ICPADS)*, Melbourne, VIC, 2015.