ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct919

Navigating Algorithmic Accountability and Ethical Governance in Autonomous Data Analytics Systems: Toward Transparent, Bias-Resistant, and Human-Centric AI Frameworks for Critical Decision-Making

Saifuddin Shaik Mohammed

Leading Healthcare Organization
California, USA
Master's in Management of Technology from New York University

Publication Date: 2025/10/28

Abstract: Intelligent data analytics systems that operate without human intervention and are powered by AI are being progressively intertwined with the decision-making processes that have a significant impact in different sectors like healthcare, finance, and criminal justice. Though these systems, in theory, make the work more efficient and insightful, their mysterious character, the possibility of algorithmic bias, and the lack of clear modes of accountability, on the other hand, expose society not only to the ethical issues but also to the social ones of considerable magnitude. This work is about the paper, which addresses the necessity of governance systems capable of regulating the situation in such a way as to ensure the responsible use of technologies, not only in terms of their development but also in terms of their deployment. I will delve deeply into the problem of algorithmic accountability from different angles, including the issue of very difficult technical audit of "black box" models and the issue of societal challenge in rectifying systemic biases embedded in training data, among other things. I come up with a full-blown, multi-layered local government model of governing TEAG, or the Tiered Ethical AI Governance Framework, combining technical instruments with the purpose of bringing about transparency and bias alleviation, together with tight procedural and organizational supervision for support. Such a human, centered approach guarantees that the self, governing systems remain compatible with ethical norms, laws, and basic human values. The integration of technical, ethical, and legal safeguards in this project reflects a shift towards the creation of AI systems that, besides being courageous and effective, are also just, clear, and answerable to the communities they exist in a deep sense.

Keywords: *Algorithmic Accountability, Ethical AI, Governance, Bias Mitigation, Explainable AI (XAI), Autonomous Systems, Data Analytics, Human-in-the-Loop, Responsible AI, Fairness Metrics, MLOps, Sociotechnical Systems*.

How to Cite: Saifuddin Shaik Mohammed (2025) Navigating Algorithmic Accountability and Ethical Governance in Autonomous Data Analytics Systems: Toward Transparent, Bias-Resistant, and Human-Centric AI Frameworks for Critical Decision-Making. *International Journal of Innovative Science and Research Technology*, 10(10), 1670-1677. https://doi.org/10.38124/ijisrt/25oct919

I. INTRODUCTION

AI and ML form the central core of a data-driven revolution, which is the most prominent feature of the 21st century. Autonomous data analytics systems, which were initially seen as only academic research, are now the main operational tools of society [1]. Algorithms in Finance are performing the estimation of creditworthiness and are detecting fraudulent activities in a snap of a second [10]. In the same manner, these algorithms in healthcare predict disease outbreaks, individualize treatment regimens, and evaluate medical imaging at a pace and extent that are beyond human capabilities [9]. The information gleaned from these methods is then fed to human experts, who make the final decisions, thus speeding up the entire process to levels not feasible by human experts alone [9]. Machine learning algorithms are used in the judiciary to determine the probability of a person reoffending. Based on this, a judge can decide whether to allow the person to go free on bail or release them on probation [5]. Such an incorporation is a huge paradigm change, as it

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct919

moves the decision-making authority from human knowledge and experience to complex computational models that claim to be objective and very efficient [2].

Instead, the rapid deployment of such technologies has largely been at odds with the drafting of moral and legal standards for governing their usage. Paradoxically, it is the very complexity of these models that endows them with power, and at the same time makes them very difficult to understand. The most advanced models, e.g., deep neural networks with millions of parameters, are even referred to as working in "black boxes."

The rationale for their decision-making is distributed in a huge network of weighted connections in such a way that human auditors cannot easily grasp it [4]. This absence of disclosure sets up a considerable obstacle to accountability. The case is, for instance, when an AI system rejects a mortgage application, suggests a wrong medical procedure, or inaccurately identifies someone as a security threat, then the most essential question of "why" is the one that we can provide a detailed answer to only in a minority of cases [5].

Moreover, the objective algorithmic nature that was expected turned out to be a trick. Artificial intelligence systems are built from data, and data reflect our past, which is a history of a very biased system and full of social injustices [6]. For example, if an AI model is given the historical hiring data of a company where mostly men were promoted, it will correlate male characteristics with success and thus female applicants will be discriminated against in a gender biased manner without even realizing it; this is exactly what happened at Amazon, which became the center of attention [7]. The algorithm is not the one doing this on purpose; it is just reflecting the biased world that it was taught to represent. Therefore, the main issue is not only a technical one but also sociotechnical. We need to find ways of building systems that are not only technically transparent but also fair from a societal point of view. Our paper recognizes these issues first and foremost and suggests a complete, end to end framework for algorithmic accountability and ethical governance that would be the guide for responsible AI innovation and its application in our most vital sectors.

II. PROBLEM STATEMENT

The central problem this research work is about, is the lack of a standardized, operable, and holistic governance framework that would regulate the operation of autonomous data analytics systems in such a way that they are accountable, transparent, and ethically aligned at all stages of their lifecycle. The absence of this governance framework creates the terrain of the so-called unmitigated risk, whose manifestation can be found in a number of critical, interconnected problem areas:

- The Accountability Vacuum: The scenario of autonomous systems leading to harm and subsequently figuring out the chain of responsibility is extremely complicated. Which one is the ultimate blame: the data scientists who created the model; the data vendors who provided biased data; the organization that used the system for a purpose it was not intended; or the end user who trusted the output? The spreading of responsibility in this case results in a legal and moral vacuum which is an area where it is almost infeasible for the displaced people to receive justice and for the companies to handle the risk properly [8]. This discrepancy not only weakens the very basis of public trust but also makes it difficult to apply the already existing legal doctrines such as tort and product liability law.
- The Interpretability-Performance Trade-Off: Most of the time, the interpretability of a model comes at the cost of its predictive accuracy, and vice versa. The most advanced models, like deep neural networks or gradient boosted trees, owe their outstanding performance to the huge complexity hiding inside, which is why their internal logic is very difficult to understand [9]. On the other hand, more straightforward and thus more interpretable models (e.g., linear regression, decision trees) are easy to understand but may not have the performance necessary for sophisticated tasks. This leaves organizations with a tough decision to make: to go with a high performing yet opaque "black box" or a transparent but less effective model. The problem is even more serious in the case of such domains as medicine or finance where both accuracy and explainability cannot be compromised.
- The Pervasiveness of Algorithmic Bias: Algorithmic bias, a complicated phenomenon, can permeate the ML pipeline at any stage. The social bias in the data can generate it used historically (e.g., discrimination based on race when granting loans) [10]; sampling bias, when the data collected is not representative of the desired population; measurement bias, where the chosen features as proxies of a target variable already have bias (e.g., using arrest rates as a proxy for crime rates); or even model bias embedded in the assumptions and objective function of the algorithm that injects and amplifies inequities [6]. Another issue that is as difficult as the definition and measurement of "fairness", is the availability of over 20 diverse mathematical definitions (e.g., Demographic Parity, Equalized Odds) that are usually contradictory to one another, thus a model cannot be "fair" per all definitions at once [11].
- Regulatory Lag and the Governance Gap: The rate at which technology changes is much faster than that of legislation. Present laws and regulations are the products of times long gone before the autonomous decision-making era, and thus leave regulators perplexed and hesitant about their next steps. Even though there are new suggestions such as the AI Act by the EU, they are only a beginning, not a full-fledged solution [12]. Without any definite legal requirements, companies have to gamble with self-regulation, and more often than not come up with broad, high level ethical principles that sound good but hardly can be transformed into specific practices in the fields of engineering and business [17].

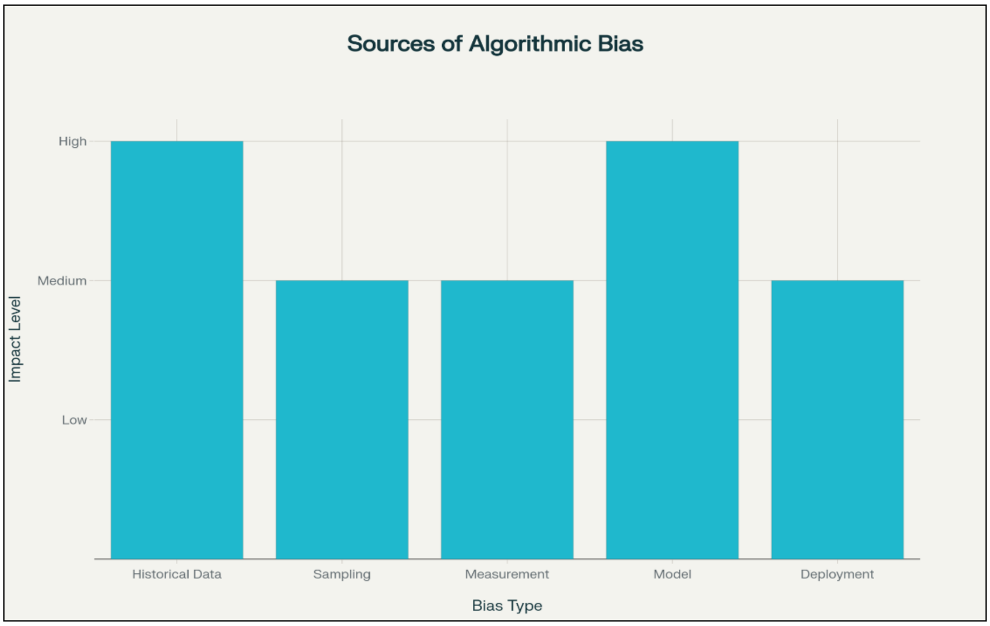

Fig 1. Sources of Algorithmic Bias and Their Prevalence

III. METHODOLOGY

In order to build a reliable and workable solution, this study applies a mixed method that combines a systematic literature review with an organized, multi-level conceptual framework development process.

➢ *Systematic Literature Review*

I have gone through the literature of articles of research, technical papers, and regulatory reports, from 2015 to the present day, in a very exhaustive manner. This was an Open Access search, meaning that it used a variety of major scholarly databases to locate the material of academic quality; among these were the IEEE Xplore, the ACM Digital Library, Scopus, and Google Scholar. Searches involved advanced search words like "AI ethics," algorithmic accountability", "explainable AI", "algorithmic bias", and "AI governance."

- Technical Tools: The study features cutting-edge XAI techniques, LIME and SHAP. The research focuses on their theoretical aspects, real-world use, and disadvantages. It also examines several biases... Pre-processing, in-processing, and post-processing bias mitigation strategies are explored in depth.

Table 1. Comparison of Explainable AI (XAI) Techniques [25] [26]

Method	Explanation	Strengths	Weaknesses
Lime	Local	Model-agnostic, interpretable	Stability issues, possible manipulation
SHAP	Local	Fair attribution, consistent	Computationally intensive for large models
Partial Dependence Plots	Global	Intuitive visualization	May miss feature interactions, less reliable for complex models

- Fairness Metrics: The study undertook comparative research of the most influential computational definitions of fairness, delving into their statistical properties and the ethical trade-offs involved in choosing one over another. For instance, Demographic Parity mandates ... of positive results is the same across all protected groups, irrespective of their underlying qualifications. On the other hand, Equal Opportunity is a less rigorous metric as it only assumes that the true positive rate is the same across groups, thus enabling equal opportunity in being correctly identified.

Table 2. Key Fairness Metrics in AI Governance [22] [23] [24]

Metric	Definition	Use Case	Limitation
Demographic Parity	Same proportion of positive outcomes for all groups.	Hiring, lending	Ignores differences in qualification rates.
Equal Opportunity	Equal true positive rate across groups.	Healthcare diagnostics	Does not address false negatives equally.
Equalized Odds	Equal true positive and false positive rates across groups.	Judicial decisions	Hard to achieve with imbalanced data.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct919

- Governance Models: The research paper undertakes a critique of existing corporate governance and risk management frameworks to derive principles and structures that can be adapted to confront the challenges posed by AI [15] [18].

➢ *Conceptual Framework Development*

The information from the literature review has been the main factor in the design of the Tiered Ethical AI Governance (TEAG) Framework. The work was done in iterations and the development was guided by the principles of systems thinking. AI was not considered as a separate technology but as a part of a larger socio-technical system. The tiered (Technical, Procedural, Organizational) structure of the framework was the reflect different layers of the system which intervention for governance effectiveness. Technical elements are based on exact computational concepts. For example, the explainability aspect uses Shapley values, a technique from cooperative game theory. This technique offers a way to provide explanations for the model's prediction by looking at the average marginal contribution of each feature for all combinations of other features thus giving the fairest way to assign influence on the outcome [13]. The main methodological contribution of the paper is the integration of such rigorous technical methods in a single, coherent governance framework.

IV. RESULTS & DISCUSSION

The systematized literature review discloses a loud and immediate agreement about the dangers of AI without governance. Still, it also shows a large division with respect to solutions. My synthesis of the findings is organized into two key themes:

➢ *Results from Literature:*

- The Limitations of Purely Technical Fixes: There is a great deal of research surrounding XAI and fairness, aware machine learning. However, the literature supports the idea of these as solutions less and less. The XAI methods can be fragile, and the interpretations of their explanations can be manipulated or misunderstood, thus giving the mistaken impression of security [16]. What is more, the inherent compromises between different fairness metrics and between fairness and accuracy imply that there cannot be one single "correct" way to remove bias from an algorithm. These are basically normative and context-dependent decisions, which require human judgment rather than automated optimization [11].
- The "Principles-to-Practice" Gap: A considerable amount of research has been devoted to the identification of high-level ethical principles for AI (e.g., transparency, justice, non-maleficence) [15]. Nevertheless, a striking difference exists between these admirable principles and their implementation within corporations and engineering workflows. Different organizations have announced their AI ethics principles. However, a look at the inside practices indicates that there are no concrete tools, training, and incentive structures to ensure that these principles are indeed followed during the development and deployment process [17].

➢ *Discussion:*

These findings, in fact, reveal an extremely significant insight: algorithmic accountability should not be considered as a problem that data scientists can solve alone. It is a sociotechnical issue that requires a big-picture approach. The presence of limitations in technical tools is like a sign that human supervision and procedural safeguards cannot be done away with. The difference between principles and practice points to the existence of formal governance organizations whose main function is to clarify the directions of responsibility and to merge the ethical considerations as the foundation of the organization's operations.

Consequently, a well-functioning governance framework should be like a bridge that connects the high-level ethical principles with the daily work of business leaders and engineers. The framework has to establish the feedback loops between technical model monitoring, procedural impact assessments, and organizational oversight at a higher level. The dialogue highlights the importance of a human in the loop (HITL) concept, which means that the automation is created to support and not substitute human judgment, especially in cases that are ambiguous, of high stakes, or exceptional [19]. This sociotechnical view is the basis for the proposed solution.

V. SOLUTION: THE TIERED ETHICAL AI GOVERNANCE (TEAG) FRAMEWORK

To solve the problems that were pinpointed and to close the gap between principles and practice, I suggest the Tiered Ethical AI Governance (TEAG) Framework. It is a detailed, multi-layered model that aims to inject ethical accountability into the essence of the AI lifecycle.

The tiers constituting the framework are three, interdependent, and mutually reinforcing:

A. Tier 1: The Technical Foundation

The lowest level is depicted by automated devices and quantitative frameworks that are straightforwardly fused into the MLOps (Machine Learning Operations) pipeline. It is the empirical data source for governance.

➢ *Continuous Bias and Fairness Auditing Module*

This is a continuous monitoring service rather than a one-time check.

- Pre-Processing: Before model training, datasets are examined for the presence of statistical biases. To make the data even, methods such as re sampling or synthetic data generation (e.g., SMOTE) may be used.
- In-Processing: The model's optimization function is directly injected with fairness constraints. That is, by changing the model's objective function it not only learns to minimize prediction error but also minimizes a penalty term that increases when a chosen fairness metric is violated. The balance between accuracy and fairness is determined by a machine setting that can be changed.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct919

- Post-Deployment Monitoring: The module keeps track of the set fairness metrics on real-time production data. For instance, it would ensure that the difference in true positive rates between demographic groups for an Equal Opportunity metric does not go beyond a small, predefined tolerance level. In case a metric exceeds this limit, it initiates an automated alert for human review.

Table 3. Overview of Bias Mitigation Strategies Across ML Pipeline [27]

Stage	Example Techniques	Key Goal
Pre-processing	Re-sampling, synthetic data, balancing	Remove or balance bias before training
In-processing	Fairness constraints in objective function	Enforce fairness during model training
Post-processing	Outcome adjustment, bias monitoring	Ensure fairness and trigger human review post-deployment

➢ *Embedded Explainability (XAI) Engine*

Every critical decision prediction must be accompanied by an explanation and logged.

- Local Explanations (e.g., LIME, SHAP): For each decision, the system produces feature attribution scores, indicating the factors that led to the specific result [4] [13].
- Global Explanations (e.g., Partial Dependence Plots): The system offers capabilities to understand the general model behavior and the effects of different features combined.
- User-Friendly Translation: There are explanations in different forms: a comprehensive one for the auditors and a simplified summary in everyday language for the person concerned.

➢ *Robustness and Security Protocols:*

- Adversarial Testing: The model is frequently challenged with adversarial attacks (e.g., minor, deliberately crafted perturbations of input data aimed at causing misclassification) to demonstrate its robustness [20].
- Privacy Preservation: For example, differential privacy is used to add statistical noise during training, thus permitting the model to learn trends in the population without saving sensitive data about individuals [20]. Decentralized approaches to privacy-preserving analytics, such as federated learning frameworks, offer additional mechanisms to protect individual privacy while enabling distributed model training across multiple data sources [28]

B. *Tier 2: The Procedural Layer*

This layer comprises human-driven procedures and standards for documentation that regulate the use of the technology.

➢ *AI Impact Assessment (AIA)*

AIA is a compulsory, auditable process that precedes any high-risk AI project. The AIA is a dynamic file that must provide answers to:

- Purpose & Context: What is the intended use and what are the potential dual-use risks?
- Data: Where does the data come from, and what are its potential biases?
- Stakeholders: Who are the people that will be affected by this system, both directly and indirectly?
- Fairness & Safety: What is the risk of unfair outcomes or errors, and is there a plan in place to mitigate them?

➢ *Structured Human-in-the-Loop (HITL) Workflows*

The framework requires the human oversight roles to be explicitly described.

- Human-in-the-Loop: For instance, a judge may actively use an AI recommendation as one piece of evidence and still make the decision (human input required for every decision).
- Human-on-the-Loop: In inspection tasks, for example, humans may evaluate model decisions and have the power to intervene if necessary (e.g., a doctor confirming an AI diagnostic suggestion).
- Human-over-the-Loop: Referring to auditing and appeals processes where people examine system performance and decide on the correctness of outcomes after the fact [19].

➢ *Comprehensive Model Governance and Documentation:*

- Model Cards: Each model deployed in production should have a “Model Card” that outlines its intended use, performance on different demographic subgroups, ethical considerations, and the limitations acknowledged [21].
- Datasheets for Datasets: The data used for the model should be described in a “Datasheet” that provides information about its collection, composition, and potential biases [6].

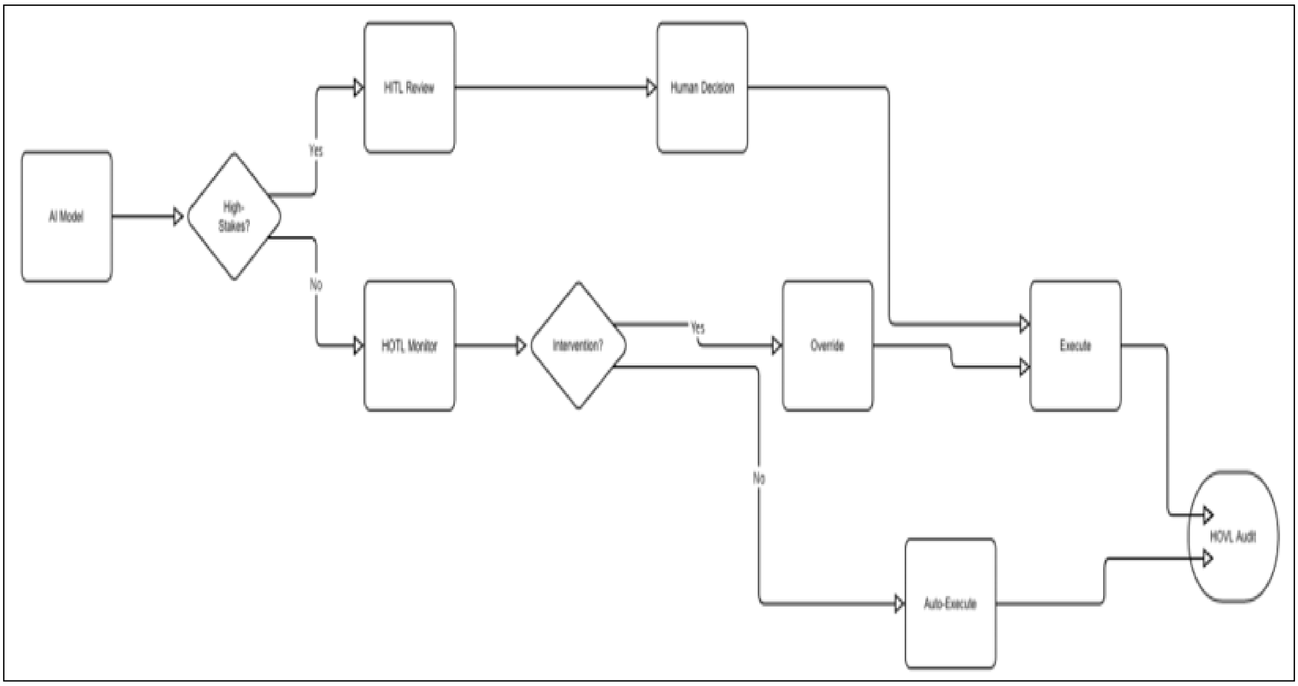

Fig 2. Human-in-the-Loop (HITL) Governance Workflow

C. Tier 3: The Organizational & Governance Oversight

The top tier is the one that puts accountability into the company's DNA, hence ensuring executive sponsorship and compliance with legal and moral standards.

➤ *AI Ethics Review Board (AERB)*

The body, which is cross-functional and permanent, has real power. Its members should not only be data scientists and lawyers but also include ethicists, social scientists, and domain experts. The AERB is the entity responsible for:

- Drafting and revising the organization's AI ethical policies.
- Giving their consent (or veto) after the examination of all high-risk AIAs.
- It is the last resort for ethical dilemmas that cannot be solved at the team level.

➤ *Clearly Defined Roles and Responsibilities*

The point of view of the Product Owner is that of the one accountable for the business case, whereas the Model Owner (a senior data scientist or engineer) is technically the one responsible for the model's performance, fairness, and monitoring. These roles report, directly or indirectly, to a C-level executive, e.g., Chief AI Officer or Chief Risk Officer.

➤ *Transparent Redress and Appeals Mechanism*

A process that is easy to access and is well-publicized for the sole purpose of enabling people to challenge decisions made by algorithms. Once an appeal is lodged, it sets off the HITL workflow (Tier 2) automatically, and the detailed XAI explanation (Tier 1) is among the pieces of evidence sent to the human reviewer. In this way, the appeal system is both effective and transparent.

VI. CONCLUSION

Self operating machines on the community fabric are offering the hopeful life of broad progress as well as the hazard of mechanized unfairness. The problems of algorithmic accountability, bias, and lack of transparency, which have been thought of as negligible side areas of technology, are the core issues that determine the future of a just and equitable digital society.

As a matter of fact, the present paper put forward the argument that confronting these issues is tantamount to giving up the idea of scattered technical solutions and the concept of ethical considerations without disclosing any details.

The suggested that the Tiered Ethical AI Governance (TEAG) Framework serves as a definite, organized, and practical instrument. In fact, through consolidating a sound technical base for observing and making comprehensible, a thorough procedural layer for evaluation and human control, and a top-level organisational structure for responsibility, the TEAG framework furnishes a comprehensive way. It is structured to realise the moral, ethical grand principles in everyday life, thereby guaranteeing that human values are still the driving force of technological innovations.

LIMITATIONS AND FUTURE WORK

Every framework misses something a few times. Implementation of the TEAG Framework needs a significant organizational commitment, resources, and a cultural shift toward ethical responsibility. The overhead can be substantial for small organizations. In addition, the success of a framework is contingent upon the thorough and good-faith performance of its activities.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct919

Research will have to continue in an effort to provide sector-specific versions of the TEAG framework, auditing protocols, and certification standards that can be used to verify compliance, as well as in the technical solutions for fairness that protect privacy and for more intuitive explainability.

In fact, the way we govern should change along with AI. Therefore, if we want to create a society in which self governing systems, despite their enormous potential, are still trustworthy colleagues with whom we can co create a world that is not only more just and efficient but also more human, it will only be possible by first introducing today accountability promoting mechanisms.

Future research should explore the integration of privacy-preserving distributed learning architectures, such as federated learning [28], within comprehensive governance frameworks like TEAG. Understanding how to maintain algorithmic accountability, transparency, and fairness in decentralized environments represents a critical frontier for responsible AI development, particularly in sensitive domains such as healthcare where data privacy and model accuracy must be simultaneously optimized.

REFERENCES

[1]. Shrestha, Y. R., & Ben-Menahem, S. M. (2019). The grand challenge of human-algorithm teaming for the future of work. Journal of Management, 45(4), 1334-1358.

[2]. Agrawal, A., Gans, J., & Goldfarb, A. (2018). Prediction Machines: The Simple Economics of Artificial Intelligence. Harvard Business Press.

[3]. McAfee, A., & Brynjolfsson, E. (2017). Machine, Platform, Crowd: Harnessing Our Digital Future. W. W. Norton & Company.

[4]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

[5]. O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown.

[6]. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for Datasets. Communications of the ACM, 64(12), 86-92.

[7]. Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters.

[8]. Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. Ethics and Information Technology, 6(3), 175-183.

[9]. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nature Machine Intelligence, 1(5), 206-215.

[10]. Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., & Walther, A. (2022). Predictably Unequal? The Effects of Machine Learning on Credit Markets. The Journal of Finance, 77(1), 5-47.

[11]. Verma, S., & Rubin, J. (2018). Fairness definitions explained. In Proceedings of the 2018 ACM/FAT Conference*.

[12]. The European Commission. (2021). Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).

[13]. Lundberg, S. M., & Lee, S. I. (2017). A Unified Approach to Interpreting Model Predictions. In Advances in Neural Information Processing Systems 30.

[14]. Hardt, M., Price, E., & Srebro, N. (2016). Equality of Opportunity in Supervised Learning. In Advances in Neural Information Processing Systems 29.

[15]. Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI. Berkman Klein Center Research Publication.

[16]. Ghassemi, M., Oakden-Abbott, J., & Ranganath, R. (2021). The false promise of armchair explanations. Nature Machine Intelligence, 3(12), 1017-1019.

[17]. Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From What to How: An Initial Review of Publicly Available AI Ethics Tools, Methods and Research to Translate Principles into Practices. Science and Engineering Ethics, 26(4), 2141-2168.

[18]. HBR Analytic Services. (2020). AI Governance and Ethics: A Call to Action. Harvard Business Review.

[19]. Shneiderman, B. (2020). Human-Centered AI: Reliable, Safe & Trustworthy. International Journal of Human–Computer Interaction, 36(6), 495-504.

[20]. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science, 9(3–4), 211-407.

[21]. Mitchell, M., Wu, S., Gessner, N., Barnes, P., Schelter, S., et al. (2019). Model Cards for Model Reporting. In Proceedings of the Conference on Fairness, Accountability, and Transparency.

[22]. Miguelañez, C. (2025, February 18). How to compare fairness metrics for model selection. Latitude Blog. https://latitude-blog.ghost.io/blog/how-to-compare-fairness-metrics-for-model-selection/

[23]. Jones, G. P., Hickey, J. M., Di Stefano, P. G., Dhanjal, C., Stoddart, L. C., & Vasileiou, V. (2020, October 8). Metrics and methods for a systematic comparison of fairness-aware machine learning algorithms. arXiv.org. https://arxiv.org/abs/2010.03986

[24]. Package ‘fairness’ july 22, 2025 title Algorithmic Fairness Metrics. (n.d.). https://cran.r-project.org/web/packages/fairness/fairness.pdf

[25]. Lcheng. (n.d.-a). https://lcheng.org/files/TAI_Fairness_chapter_2022.pdf

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct919

[26]. Castelnovo, A., Crupi, R., Greco, G., Regoli, D., Penco, I. G., & Cosentini, A. C. (2022, March 10). A clarification of the nuances in the fairness metrics landscape. Nature News. https://www.nature.com/articles/s41598-022-07939-1
[27]. Arxiv. (n.d.-a). https://arxiv.org/pdf/2010.03986.pdf
[28]. Saifuddin Shaik Mohammed. "A Decentralized Approach to Privacy-Preserving Data Analysis using Federated Learning." Volume. 10 Issue.9, September-2025 International Journal of Innovative Science and Research Technology (IJISRT),2091-2096 https://doi.org/10.38124/ijisrt/25sep1145