



“ALGORITHMIC BIAS AND THE LIMITS OF U.S. LAW: LESSONS FROM THE EU AI ACT’S RISK-BASED FRAMEWORK”

Eshmurodov Zoyir.

Head and Chief Lawyer(CLO) of the law firm “ZOXA
LAW”, established in the United States of America

Abstract *This article examines the growing problem of algorithmic bias in the United States and argues that the current sector-specific legal framework is insufficient to address the discriminatory harms produced by AI systems. It highlights how automated tools used in employment, lending, housing, healthcare, and criminal justice often replicate or intensify social inequalities, while U.S. laws such as Title VII, ECOA, and the Fair Housing Act provide only reactive, post-hoc remedies. The article contrasts this fragmented model with the European Union’s AI Act, which introduces a comprehensive, risk-based regulatory structure requiring pre-deployment audits, transparency, human oversight, and bias testing for high-risk AI systems. By comparing the reactive U.S. approach with the EU’s preventive model, the article concludes that the United States could benefit from adopting hybrid reforms that integrate risk assessments and mandatory bias mitigation into existing civil-rights statutes. Such an approach would strengthen accountability, reduce algorithmic discrimination, and create a more future-proof legal environment for the governance of automated decision-making technologies.*

Keywords *algorithmic bias, artificial intelligence, machine learning, automated decision-making, fairness in AI, discrimination law,*

civil rights, equal protection, data privacy law, bias in algorithms, predictive analytics, transparency and accountability, GDPR, CCPA, employment law and AI, criminal justice algorithms, risk assessment tools, legal liability for AI, ethical AI, technology and human rights

Introduction

Artificial intelligence now influences decisions in hiring, lending, housing, education, and policing, but these automated systems often produce discriminatory outcomes. In the United States, studies show that certain AI hiring tools reject up to 35% more women than men, and algorithmic mortgage systems deny Black applicants at rates 40–80% higher than similarly qualified white applicants. These disparities reveal a major weakness in the U.S. legal framework: regulation is sector-based, meaning each industry is governed by its own separate laws—such as Title VII for employment, the Equal Credit Opportunity Act for lending, and the Fair Housing Act for housing. None of these statutes were originally designed for machine learning, and enforcement is usually reactive, occurring only after discrimination is discovered.

By contrast, the EU AI Act¹, adopted in 2024, uses a risk-based model that categorizes AI systems into four levels—unacceptable risk, high risk, limited risk, and minimal risk. High-risk systems, such as AI used in hiring or credit scoring, must undergo strict compliance checks, including transparency, human oversight, and bias testing before deployment. This approach aims to detect discrimination early rather than waiting for legal violations to occur. As U.S. regulators debate potential federal AI legislation, the EU's preventive structure provides a

¹ <https://artificialintelligenceact.eu/>

roadmap for reducing algorithmic harms. Understanding the contrast between these two models is essential for designing future-proof protections that promote both innovation and fairness in automated decision-making.

The Problem: Automated Discrimination in the U.S.

Automated decision-making systems are now used across almost every major sector in the United States—employment, housing, healthcare, education, policing, immigration, and consumer finance. While these technologies promise efficiency, they often replicate or even amplify existing inequalities. The core challenge is that the United States does not have a single federal statute specifically regulating algorithmic discrimination. Instead, the system relies on older civil-rights laws, each limited to a particular domain. Title VII covers employment discrimination², the Equal Credit Opportunity Act (ECOA)³ covers lending, the Fair Housing Act (FHA) governs housing, and HIPAA and various state laws address certain health-related issues. These statutes were drafted in the 1960s–1990s and were never designed for opaque machine-learning models that operate at scale and are capable of producing millions of decisions per day.

Because of this structure, the U.S. enforcement model is largely reactive, not preventive. Regulators like the EEOC, CFPB, and HUD typically step in only after an individual files a complaint—which is difficult when people never learn how the algorithm reached its decision. Studies illustrate how pervasive these hidden harms can be: an influential 2019 analysis found that a healthcare risk algorithm used for

² <https://www.eeoc.gov/statutes/title-vii-civil-rights-act-1964>

³ <https://www.justice.gov/crt/equal-credit-opportunity-act-3>

over 200 million patients systematically underestimated the needs of Black patients, reducing their access to care. Research at MIT showed⁴ that commercial facial-recognition systems misidentified darker-skinned women at rates above 30%, compared to error rates below 1% for lighter-skinned men. Investigations in the housing and lending sectors reveal similar patterns, with algorithmic mortgage systems denying Black applicants 40–80% more often, even when they have comparable qualifications.

Because the U.S. does not require companies to conduct pre-deployment testing, documentation, or transparency reporting, many discriminatory systems operate for years before they are discovered. This lack of mandatory oversight allows algorithmic bias to spread quietly through essential sectors of daily life, creating systemic harms that disproportionately affect already vulnerable communities. The result is a regulatory landscape where the law protects individuals only after damage has been done—highlighting the limits of the U.S. sectoral approach in an age of increasingly autonomous and data-driven technologies.

EU AI Act: A Preventive, Risk-Based Strategy

The EU AI Act, adopted in 2024, is the world's first comprehensive legal framework for artificial intelligence, and it takes a preventive, risk-based approach that is fundamentally different from the U.S. model. Instead of regulating sectors one by one, the Act classifies AI systems into four risk tiers:

⁴ <https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212>

1. Unacceptable Risk – systems that pose a clear threat to safety or fundamental rights, such as social scoring or manipulative biometric systems. These systems are completely banned.
2. High Risk – AI used in areas like hiring, credit scoring, education, border control, and healthcare. These systems must undergo strict pre-deployment obligations, including documentation, accuracy testing, cybersecurity measures, human oversight, and bias audits.
3. Limited Risk – systems such as chatbots or recommendation algorithms that require transparency so users know they are interacting with AI.
4. Minimal Risk – most consumer applications, like video-game AI or spam filters, which face no additional rules.

This tiered structure forces companies to evaluate risks before deployment, which is why the EU model detects discrimination earlier and more reliably. Under the Act, a company cannot deploy a high-risk algorithm—such as an automated hiring tool or credit-risk model—without first proving that the system is accurate, explainable, and non-discriminatory. Regulators can demand technical documentation, training data summaries, and audit reports before the system is used on the public, something U.S. law currently does not require. Early studies by the European Commission estimate that mandatory risk assessments could reduce algorithmic errors and discriminatory outcomes by 20–30% in high-risk sectors within the first years of implementation⁵.

Overall, the EU's approach is proactive: it aims to prevent harm before it happens, while the U.S.'s sectoral laws intervene after the harm is discovered. This distinction explains why many policymakers view the

⁵ <https://rm.coe.int/discrimination-artificial-intelligence-and-algorithmic-decision-making/1680925d73>

EU AI Act as a more future-proof model for governing automated decision-making.

Where the U.S. Falls

The United States currently lacks a comprehensive legal framework for regulating AI, which leaves significant gaps in protecting against automated discrimination.

- No general AI law: Unlike the EU, which has a single AI Act, U.S. law relies on sector-specific statutes—such as Title VII for employment, the Equal Credit Opportunity Act for lending, and the Fair Housing Act for housing—without overarching rules for algorithmic systems.
- No mandatory risk assessments: Companies are not required to evaluate AI systems for bias or potential harm before deployment, meaning discriminatory outcomes can persist unnoticed for years.
- No dataset quality requirements: There is currently no legal obligation to ensure that training data is representative, unbiased, or accurate, even though studies show that biased datasets can lead to up to 35% higher rejection rates for women in hiring algorithms and 40–80% higher mortgage denial rates for Black applicants.
- No transparency mandates: Businesses do not have to disclose how AI systems make decisions, leaving consumers and regulators in the dark about potential harms.
- Enforcement is slow and reactive: Federal agencies like the EEOC, CFPB, and HUD act only after complaints are filed, making it difficult to prevent widespread discrimination. Collectively, these gaps illustrate why the U.S. approach struggles to address algorithmic bias proactively, leaving both vulnerable populations and regulators at a significant disadvantage compared to the EU's preventive, risk-based model.

Conclusion

The contrast between the United States' sectoral, reactive approach and the European Union's preventive, risk-based AI framework highlights both the strengths and limitations of current regulatory strategies. The sector-specific laws, such as Title VII, ECOA, and the Fair Housing Act, are well-established and allow for targeted enforcement within specific industries. These statutes do not address AI holistically, offer no pre-deployment bias testing, and rely heavily on individuals to detect harm, leaving many discriminatory outcomes invisible. The EU AI Act proactively categorizes systems by risk, mandates pre-deployment audits, bias testing, transparency, and human oversight, reducing the likelihood of algorithmic harm before it occurs. Some critics argue that the system could impose compliance costs that slow innovation, particularly for small and medium-sized enterprises. A hybrid approach could combine the U.S. focus on targeted enforcement with the EU's preventive measures. For example, the U.S. could retain its sectoral laws but require risk assessments, documentation, and bias testing for high-impact AI systems, such as hiring algorithms, credit scoring models, and predictive policing tools. Such a framework would protect consumers and vulnerable populations while preserving flexibility and innovation. Ultimately, learning from the EU's risk-based model could allow the United States to develop a forward-looking regulatory strategy that balances fairness, accountability, and technological progress.

Reference List

1. Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning: Limitations and opportunities*. <https://fairmlbook.org/>
2. O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown Publishing.

3. Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
4. Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10(1), 113–174. <https://doi.org/10.1093/jla/lax001>
5. Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Proceedings of Machine Learning Research*, 81, 149–159.
6. European Union. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation)*. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>
7. U.S. Equal Employment Opportunity Commission. (2023). *EEOC guidance on AI and employment discrimination*. <https://www.eeoc.gov/ai-guidance>
8. Executive Office of the President. (2016). *Big data: A report on algorithmic fairness and civil rights*. Washington, DC: U.S. Government. https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf
9. State of California. (2022). *California Consumer Privacy Act (CCPA)*. <https://oag.ca.gov/privacy/ccpa>
10. Loomis v. Wisconsin, 881 N.W.2d 749 (Wis. 2016).
11. State v. Jackson, 388 P.3d 935 (Wash. 2017).
12. United States v. Microsoft Corp., 138 S. Ct. 1186 (2018) (algorithmic/data-driven issues in context of privacy and compliance).
13. AI Now Institute. (2018). *Discriminating systems: Gender, race and power in AI*. New York University. <https://ainowinstitute.org/discriminatingystems.pdf>
14. Partnership on AI. (2020). *Fair, transparent, and accountable AI*. <https://www.partnershiponai.org/fairness/>