

Algorithmic Bias and Equal Protection: Rethinking Anti-Discrimination Law for the Age of Artificial Intelligence

1. Dr. Rehan Zafar, Faculty of Law, Bahria University, Islamabad — rehan.zafar@gmail.com
2. Mehak Tariq, Department of Law, University of Okara — mehak.tariq.law@gmail.com

Abstract

As artificial intelligence (AI) systems increasingly mediate access to credit, employment, education, and criminal justice, algorithmic bias has emerged as a constitutional question of first order. Traditional equal protection jurisprudence—built upon intentional discrimination and suspect classifications—struggles to address structural harms generated by opaque and data-driven decision processes. This article examines how the Fourteenth Amendment’s Equal Protection Clause can be re-interpreted for an algorithmic state. It traces doctrinal limits of intent-based discrimination tests, evaluates the constitutional treatment of disparate-impact analysis, and explores the normative potential for algorithmic accountability within the existing constitutional framework. Drawing upon U.S. Supreme Court precedent, interdisciplinary scholarship, and comparative lessons from data-governance theory, the article argues for a reconstructed constitutional paradigm that embeds transparency, fairness, and algorithmic due process within equal protection analysis. Ultimately, it calls for the evolution of constitutional interpretation that recognizes algorithmic systems as potential instruments of state action demanding heightened scrutiny and procedural safeguards.

SUBMISSION DATE: OCTOBER 2, 2023
ACCEPTANCE DATE: OCTOBER 8, 2023
PUBLICATION DATE: OCTOBER 10, 2023

Keywords

Artificial Intelligence; Algorithmic Bias; Equal Protection; Constitutional Law; Disparate Impact; Due Process; Algorithmic Governance; Data Ethics; AI Regulation; Fourteenth Amendment; Civil Rights; Machine Learning Fairness

Introduction

The Fourteenth Amendment’s Equal Protection Clause embodies the constitutional promise that “no state shall ... deny to any person within its jurisdiction the equal protection of the laws.” Designed to dismantle state-sanctioned racial hierarchies, this provision has historically anchored American anti-discrimination law. Yet its

jurisprudence remains grounded in the detection of human intent. The Supreme Court’s canonical decisions—*Washington v. Davis*, *Personnel Administrator v. Feeney*, and *Arlington Heights v. Metropolitan Housing Corp.*—define discrimination as purposeful unequal treatment rather than disparate outcomes. This human-centric conception now falters in the algorithmic age, where bias can emerge from data patterns and model architectures rather than animus [2], [6].

Artificial-intelligence systems replicate and amplify existing inequalities because they learn from historically skewed data sets [5], [14]. Algorithms employed in hiring, credit scoring, predictive policing, and sentencing often encode correlations that correlate with race, gender, or socioeconomic status. The resulting discrimination is structural, statistical, and frequently unintentional [7], [11]. Courts, however, lack doctrinal tools to recognize such harm as constitutionally cognizable. The constitutional concept of intent—as articulated in *Davis* and *Feeney*—demands proof of subjective purpose, which machine systems by definition lack [17].

The central problem, therefore, is epistemic: equal-protection analysis presupposes human deliberation, while algorithmic governance replaces judgment with optimization. The challenge is not merely technological but constitutional—whether equality jurisprudence can evolve to regulate automated decision-making that has no “mind” to discriminate yet produces outcomes indistinguishable from prejudice. Scholars like Barocas and Selbst call this the crisis of “disparate impact without intent” [2]. Others, including Kramer [22] and Freeman [40], argue that constitutional law must recognize algorithmic bias as a form of structural inequality akin to systemic discrimination in housing or education.

Building upon this discourse, the present article reconsiders the architecture of equal-protection doctrine through three analytical moves. First, it reconstructs the conceptual relationship between state action and algorithmic decision-making, demonstrating how governmental delegation to machine systems constitutes constitutionally reviewable conduct. Second, it examines doctrinal strains within intent-based jurisprudence and the marginalization of disparate-impact reasoning in constitutional law. Third, it proposes a policy-oriented framework that integrates algorithmic accountability, procedural fairness, and public oversight into the constitutional analysis of equality.

In doing so, it contributes to the growing scholarship that links constitutional values to technological governance [4], [9], [13], [19]. The argument proceeds from a normative conviction that the Equal Protection Clause must not remain static while power migrates from officials to algorithms. Constitutionalism, as both restraint and empowerment, must adapt to the digital state where decisions are mediated through code.

I. Constitutional Foundations of Equal Protection

Equal protection doctrine has long oscillated between two conceptions of equality: formal and substantive. Formal equality demands identical treatment irrespective of group membership; substantive equality seeks to remedy structural disadvantage. The Supreme Court’s jurisprudence has largely favored the former, applying strict scrutiny only to classifications that reflect a history of invidious discrimination such as race or national origin, and rational-basis review to most others [52], [53].

This binary structure presupposes identifiable classes and conscious state decisions. Yet algorithmic systems operate through proxies rather than classifications: zip codes, purchasing history, or linguistic cues that statistically align with race or gender [2], [5]. In *Washington v. Davis* (1976), the Court held that a facially neutral employment test with racially disparate outcomes did not violate equal protection absent proof of intent. That reasoning—anchored in the psychological model of discrimination—cannot account for algorithmic correlation, where discriminatory outcomes emerge from statistical learning rather than purpose.

The consequence is a constitutional blind spot. If an AI-based parole model disproportionately denies release to Black inmates because historical data associate recidivism with race, such bias may escape constitutional scrutiny unless plaintiffs can show that programmers or officials acted with discriminatory motive [6], [17]. As Lee observes, “the current doctrine makes algorithmic discrimination invisible to the Constitution” [6].

A second foundational element is the **state-action doctrine**. Equal protection binds only governmental conduct, not private actors. Yet in an algorithmic ecosystem where states purchase, license, or rely on proprietary models developed by private vendors, the boundary between public and private power dissolves. When an administrative agency adopts a risk-assessment tool to allocate welfare benefits or determine sentencing, its reliance on algorithmic output constitutes state action because the system’s decisions are effectively those of the state [22], [42]. Recognizing such delegation as constitutional conduct is imperative if equal protection is to retain relevance.

Moreover, as A. Raza notes, equality “is not a self-executing concept but a living instrument requiring interpretive renewal” [4]. His observation about the evolutionary character of equality jurisprudence underscores the need for interpretive elasticity in the face of technological transformation. Constitutional fidelity thus requires doctrinal evolution, not inertia.

II. The Problem of Intent in an Algorithmic State

Intent-based analysis remains the most significant doctrinal obstacle to recognizing algorithmic bias as unconstitutional. Under *Davis* and *Feeney*, a law or policy violates equal protection only when a discriminatory purpose “was a motivating factor.” This requirement of subjective intent fits poorly with statistical or structural discrimination [24].

Algorithms lack minds, motives, or consciousness. Their “intent” manifests in design choices, data selection, and optimization objectives [7], [16]. Bias may result from training data that reflect existing inequalities, from proxy variables that encode race, or from performance metrics that prioritize accuracy over fairness [27], [38]. None of these involve intent as the Court understands it. As Barocas and Selbst note, “discrimination by data” operates through correlation, not cognition [2].

The **doctrinal paradox** is that the Constitution’s most powerful equality guarantee cannot reach the most pervasive contemporary forms of inequality. While statutory regimes such as Title VII permit disparate-impact claims, the Court in *Washington v.*

Davis explicitly rejected that approach for constitutional claims, fearing it would transform the judiciary into a “super-legislature” reviewing all racial disparities. Yet in the algorithmic context, refusing to acknowledge disparate impact amounts to judicial abdication.

Kramer’s analysis of “Equal Protection by Design” suggests that intent doctrine can evolve by recognizing foreseeability and indifference as proxies for purpose [22]. When policymakers knowingly deploy opaque algorithms despite evidence of bias, their willful blindness should satisfy the intent requirement. This reasoning parallels *Farmer v. Brennan*’s deliberate-indifference standard in Eighth-Amendment jurisprudence. Applying similar logic to equal protection would capture institutional negligence in algorithmic deployment.

A comparable argument arises in administrative-law scholarship. Ho and Bacon propose that “algorithmic fairness” obligations within agencies function as constitutional due process analogues [10]. Embedding these duties within equal-protection analysis would allow courts to assess whether governmental reliance on biased systems constitutes arbitrary or discriminatory state action.

Finally, intent doctrine’s rigidity undermines the preventive function of constitutional law. As M. Eubanks documents, welfare algorithms that profile the poor operate as structural discrimination even absent malicious intent [17]. A constitution that demands proof of prejudice before recognizing injustice fails its moral purpose. In the algorithmic age, the constitutional inquiry must shift from *why* discrimination occurs to *how* it occurs and *whether* the state could have prevented it.

III. Disparate Impact and the Limits of Constitutional Formalism

The Court’s reluctance to constitutionalize disparate-impact theory reflects a commitment to formal equality. Yet formal equality presumes that neutral rules yield neutral outcomes, a presumption belied by machine learning [40], [43]. Disparate-impact doctrine emerged statutorily under Title VII and the Fair Housing Act, where Congress recognized that systemic biases can persist without intent. In *Griggs v. Duke Power Co.* (1971), the Court upheld this reasoning in a statutory context but never extended it constitutionally.

AI systems compel reconsideration. Algorithms routinely generate disparate outcomes despite neutral design [11], [45]. Empirical studies show that predictive models in health care and credit scoring systematically disadvantage minorities because they optimize for cost efficiency or repayment probability rather than fairness [45]. Equal protection’s failure to engage these structural disparities risks rendering the Constitution obsolete as a tool of justice.

E. Kim argues that the Fourteenth Amendment’s core purpose—eradicating caste-like structures—justifies importing disparate-impact reasoning into constitutional analysis [26]. Freeman similarly contends that fairness metrics in AI represent modern expressions of equal protection’s substantive spirit [40]. This perspective aligns with Fuller’s “inner morality of law,” which demands congruence between legal rules and their social consequences [53].

Reformulating disparate impact as a constitutional doctrine would not require abandoning intent; rather, it would supplement it. Courts could recognize a rebuttable presumption of unconstitutionality when algorithmic systems produce statistically significant disparities along protected lines, shifting the burden to the state to justify the model's necessity and the unavailability of less discriminatory alternatives. Such a framework mirrors strict-scrutiny analysis without conflating correlation with culpability.

Pragmatically, this approach aligns with emerging policy initiatives advocating algorithmic-impact assessments and fairness audits within public agencies [10], [42]. It also resonates with A. Raza et al.'s argument that "automation in judicial administration demands evaluative frameworks that reconcile efficiency with accountability" [18]. By analogy, algorithmic governance in other state functions must reconcile optimization with constitutional equality.

Ultimately, constitutionalizing disparate-impact reasoning would reaffirm the Fourteenth Amendment's normative foundation: equality as fairness in effect, not merely in form. As Hildebrandt observes, law's legitimacy in technological societies depends on its capacity to internalize the logics of code without surrendering its moral core [13].

IV. State Action and the Constitutional Accountability of Algorithms

The state action doctrine serves as the constitutional hinge upon which equal protection applies. The Fourteenth Amendment restrains only governmental conduct, but modern governance often occurs through privatized technologies. Algorithms used in welfare eligibility, parole assessments, and credit-scoring systems may be designed and maintained by private vendors, yet they perform quintessentially public functions when adopted by agencies or courts [42], [47]. This diffusion of responsibility threatens to create "accountability gaps" where no entity can be held constitutionally liable for discriminatory outcomes.

Constitutional theory offers two ways to resolve this. First, the **public function doctrine** extends state action to private entities performing functions "traditionally and exclusively reserved to the state." Predictive policing algorithms, for example, exercise coercive authority akin to law enforcement, thus satisfying this criterion [44]. Second, the **entanglement test** treats private conduct as state action when the government is significantly involved in or benefits from it. When agencies purchase or rely upon predictive models that systematically disadvantage protected groups, their adoption constitutes endorsement of the model's discriminatory logic [21].

As A. Raza and co-authors observe, automation within judicial or administrative systems cannot escape accountability simply by technological delegation [18]. The Constitution binds not the machinery of governance but the governance of machinery. If the state knowingly or recklessly employs models that produce disparate racial or gender outcomes, it acts "under color of law" in the constitutional sense. This position accords with Newman's argument that algorithmic governance blurs the line between private autonomy and public power [33].

Yet judicial recognition of algorithmic state action remains limited. Courts often defer to contractual or proprietary boundaries, accepting the state’s claim that private vendors’ algorithms lie beyond constitutional review [47]. This approach risks immunizing discrimination under the veneer of innovation. To restore constitutional coherence, courts must reinterpret state action functionally: when algorithmic decisions materially determine the distribution of public benefits or burdens, they constitute state action, regardless of formal ownership.

Such a move would reaffirm the rule articulated in *Shelley v. Kraemer* (1948), where judicial enforcement of private covenants was deemed state action. Likewise, government endorsement or enforcement of algorithmic outputs embeds them within constitutional responsibility. The Fourteenth Amendment’s normative reach must extend to the architecture through which the modern state governs.

V. Algorithmic Rationality and Standards of Review

Equal protection doctrine employs a tiered standard of review: strict scrutiny for suspect classifications, intermediate scrutiny for quasi-suspect classes, and rational basis for all others. This structure, designed for human decision-making, is ill-suited for algorithmic governance [22]. When algorithms produce racial or gender disparities, courts struggle to identify the appropriate standard: is the bias intentional, systemic, or incidental? The result is doctrinal paralysis.

Strict scrutiny requires that a classification be narrowly tailored to achieve a compelling governmental interest. Yet algorithms rarely “classify” explicitly; they correlate. The opacity of machine learning obscures whether a given variable operates as a racial proxy or a neutral predictor [11], [29]. Rational-basis review, conversely, is too deferential to capture structural harms. The challenge, therefore, is to craft an **algorithmic standard of review** that preserves constitutional principles while acknowledging technological complexity.

Scholars propose several models. Kramer advocates an “enhanced rationality review” that evaluates the proportionality between predictive accuracy and disparate impact [22]. Freeman [40] similarly argues for a “technological strict scrutiny” where government use of AI must satisfy necessity, transparency, and less-discriminatory-alternative tests. These approaches echo Fuller’s procedural morality, demanding coherence between constitutional ends and technological means [53].

Integrating such tests into equal protection analysis could harmonize constitutional review with administrative accountability. Agencies deploying AI should bear the burden of demonstrating that models were validated for bias, audited for fairness, and subject to public explanation [10], [42]. Failure to meet these conditions should render the deployment unconstitutional, regardless of intent. This shifts judicial scrutiny from subjective purpose to objective justification—a move consonant with both rule-of-law theory and contemporary understandings of algorithmic governance [13], [38].

Moreover, as A. Raza et al. suggest, when judicial or administrative automation is introduced without adequate procedural safeguards, it undermines the legitimacy of the justice system itself [18]. Equal protection, therefore, must evolve from a reactive doctrine of culpability to a proactive principle of institutional design. Courts should

evaluate not only whether discrimination occurred, but whether the state implemented reasonable mechanisms to prevent it.

VI. Data, Privacy, and the Substantive Dimension of Equality

Algorithmic discrimination does not occur in isolation from informational asymmetry. The data used to train machine-learning models encode social hierarchies, and the absence of transparency prevents individuals from contesting or understanding adverse outcomes [14], [25]. Equal protection and privacy thus converge: both safeguard personal dignity against arbitrary state power.

A. Raza and colleagues underscore this connection in their study of privacy law's evolution, arguing that technological governance must align with constitutional principles of fairness and autonomy [14]. This insight resonates with Warren and Brandeis's foundational concept of privacy as the "right to be let alone" [48]—a right that, in the digital context, translates into protection from data profiling and surveillance-based discrimination.

Data governance and equality share a structural logic. Both demand procedural transparency, justification, and oversight. When individuals are subject to algorithmic decisions without explanation or recourse, they are denied both due process and equal protection [6], [47]. Transparency is thus a constitutional value, not merely an ethical recommendation.

To operationalize this, policymakers must embed **algorithmic transparency requirements** within constitutional frameworks. Just as administrative law mandates reason-giving for agency actions, equal protection should require disclosure of model rationale and impact assessments. The failure to disclose relevant data features or fairness metrics should constitute a constitutional defect in itself.

This integration of privacy and equality further underscores the normative unity of constitutional rights in the algorithmic era. The Constitution's purpose, as Sandel and Dworkin suggest, is not formal symmetry but moral coherence: the consistent recognition of human dignity across evolving forms of power [51], [52]. In algorithmic governance, dignity depends on visibility—the ability to know, challenge, and influence the systems that govern one's life. Equal protection must therefore secure informational parity as a precondition for equality.

VII. Toward a Policy Framework for Algorithmic Equality

Doctrinal reform alone cannot safeguard equality in algorithmic governance. The constitutional project must be supplemented by policy innovation that embeds fairness and accountability into every layer of the technological ecosystem. This section proposes a **three-tiered policy framework** grounded in constitutional principles.

A. Procedural Integration

Courts and agencies should require **Algorithmic Impact Statements (AIS)** analogous to environmental impact assessments. Before adopting AI tools that affect rights or benefits, the government must analyze potential disparate impacts, publish

methodologies, and invite public comment [10], [42]. Such procedures operationalize equal protection's deliberative function, ensuring that algorithmic governance remains transparent and contestable.

Additionally, judicial review should include algorithmic evidence standards. Following Kroll et al.'s concept of "accountable algorithms," litigants must have access to explanations sufficient to test constitutionality [46]. This proceduralization of equality would extend due process principles into the algorithmic domain.

B. Substantive Standards

The Constitution should recognize a *duty of algorithmic justification*: whenever government action relies on predictive models, the state must demonstrate that its reliance is empirically valid, necessary, and proportionate. This aligns with the proportionality test used in other constitutional systems and parallels strict-scrutiny reasoning [22], [40]. Governmental use of biased or unvalidated algorithms should be presumptively unconstitutional absent compelling necessity.

C. Institutional Accountability

Finally, the policy framework must institutionalize algorithmic oversight through independent auditing bodies or ombudspersons. These entities should have statutory authority to inspect government-deployed models, assess fairness, and publish findings. Judicial reliance on private contractors does not absolve constitutional responsibility; oversight mechanisms must extend to third-party vendors engaged in state action [21], [33].

In sum, the policy imperative is to constitutionalize AI governance—to embed equality and fairness within technical design and administrative practice alike. The Equal Protection Clause should function not merely as a reactive safeguard but as an affirmative principle guiding the architecture of the digital state.

Conclusion

The challenge of algorithmic bias reveals the fragility of a constitutional doctrine rooted in human intent. As machine systems mediate public power, formal equality collapses under the weight of structural data inequalities. The Fourteenth Amendment's enduring promise—equal protection of the laws—demands reinterpretation to confront this new reality.

This article has argued that constitutional law must evolve along three dimensions. First, redefining state action to encompass algorithmic governance ensures that the Constitution binds not only human actors but the digital architectures through which they rule. Second, revising intent doctrine to include foreseeability, indifference, and negligent adoption bridges the gap between human culpability and algorithmic harm. Third, incorporating disparate-impact analysis within constitutional reasoning reclaims the Amendment's substantive purpose: to dismantle structural hierarchies, however mechanized their form.

Policy measures such as Algorithmic Impact Statements, fairness audits, and independent oversight bodies can translate these principles into practice. By constitutionalizing algorithmic governance, the state reaffirms its moral commitment to equality as fairness in outcome, not mere neutrality in code.

As A. Raza aptly observed, equality is “a living instrument requiring interpretive renewal” [4]. That renewal is now constitutional necessity. The republic must not allow automation to become the new arbiter of human worth. The Constitution must remain the measure of justice—even when written in code.

References

- [1] F. Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information*, Harvard University Press, 2015.
- [2] S. Barocas and A. D. Selbst, “Big Data’s Disparate Impact,” *California Law Review*, vol. 104, pp. 671–732, 2016.
- [3] J. Citron and F. Pasquale, “The Scored Society: Due Process for Automated Predictions,” *Washington Law Review*, vol. 89, pp. 1–33, 2014.
- [4] A. Raza, “Equality before Law and Equal Protection of Law: Contextualizing its Evolution in Pakistan,” *Pakistan Law Journal*, 2023.
- [5] C. O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, Crown, 2016.
- [6] R. C. Lee, “Algorithmic Due Process and the Constitution,” *Stanford Technology Law Review*, vol. 22, pp. 287–331, 2019.
- [7] S. Barocas, M. Hardt, and A. Narayanan, *Fairness and Machine Learning: Limitations and Opportunities*, MIT Press, 2023.
- [8] D. M. West, *The Future of Work: Robots, AI, and Automation*, Brookings Institution Press, 2018.
- [9] A. Raza, M. A. Chohan, N. Khan, G. Ali, and N. A. Tayyab, “Artificial Intelligence and Criminal Liability: Rethinking Criminal Liability in the Era of Automated Decision Making,” 2023.
- [10] D. E. Ho and D. E. Bacon, “Algorithmic Fairness in the Administrative State,” *Harvard Law Review Forum*, vol. 133, pp. 171–208, 2020.
- [11] J. Angwin et al., “Machine Bias,” *ProPublica*, May 23, 2016.
- [12] K. P. Donohue, “The Limits of Anti-Discrimination Law in the Age of Algorithms,” *Yale Law Journal Forum*, vol. 130, pp. 88–109, 2020.
- [13] L. Hildebrandt, *Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology*, Edward Elgar, 2015.
- [14] A. Raza, A. Yasin, S. Khalid, S. B. R. Naqvi, and U. Noreen, “From Bytes to Boundaries: Finding the Fate of Privacy Law in the Era of Technology,” 2023.
- [15] D. Kaye, “Artificial Intelligence and Human Rights: Report to the UN General Assembly,” United Nations, 2018.
- [16] J. Balkin, “The Three Laws of Robotics in the Age of Big Data,” *Ohio State Law Journal*, vol. 78, pp. 1217–1241, 2017.
- [17] M. Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*, St. Martin’s Press, 2018.
- [18] B. Munir, A. Raza, S. Khalid, and S. M. Kasuri, “Automation in Judicial Administration: Evaluating the Role of Artificial Intelligence,” 2023.
- [19] P. T. Kim, “Data-Driven Discrimination at Work,” *William & Mary Law Review*, vol. 58, pp. 857–936, 2017.

- [20] J. O. McGinnis and R. P. Pearce, “The Great Disruption: How Machine Intelligence Will Transform the Role of Lawyers,” *Fordham Law Review*, vol. 82, pp. 3041–3066, 2014.
- [21] R. Calo, “Artificial Intelligence Policy: A Primer and Roadmap,” *UC Davis Law Review*, vol. 51, pp. 399–435, 2017.
- [22] R. B. Kramer, “Equal Protection by Design: Algorithmic Bias and the Fourteenth Amendment,” *Georgetown Law Journal*, vol. 111, pp. 245–293, 2022.
- [23] K. Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*, Yale University Press, 2021.
- [24] D. Citron, *Hate Crimes in Cyberspace*, Harvard University Press, 2014.
- [25] R. Zarsky, “Understanding Discrimination in the Scored Society,” *Washington Law Review*, vol. 89, pp. 1375–1412, 2014.
- [26] E. Kim, “Equal Protection in the Age of AI: Revisiting Disparate Impact,” *Michigan Journal of Law & Technology*, vol. 27, pp. 117–160, 2021.
- [27] S. Binns, “Fairness in Machine Learning: Lessons from Political Philosophy,” *Conference on Fairness, Accountability, and Transparency (FAT)*, 2018.
- [28] S. Jasanoff, “The Ethics of Invention: Technology and the Human Future,” W. W. Norton, 2016.
- [29] D. Crawford and M. Paglen, “Excavating AI: The Politics of Images in Machine Learning Training Sets,” *International Journal of Cultural Studies*, vol. 22, pp. 1–16, 2019.
- [30] M. L. D. Seligman, “Algorithmic Justice and Disparate Impact,” *New York University Law Review*, vol. 95, pp. 141–189, 2020.
- [31] R. C. Post, *Democracy, Expertise, and Academic Freedom: A First Amendment Jurisprudence for the Modern State*, Yale University Press, 2012.
- [32] J. Tasioulas, “The Rule of Law in the Age of AI,” *Oxford Legal Studies Research Paper*, no. 11, 2021.
- [33] D. G. Newman, “Algorithmic Governance and the Limits of Equality,” *University of Chicago Law Review Online*, vol. 88, pp. 1–17, 2021.
- [34] R. Abebe et al., “Roles for Computing in Social Change,” *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, 2020.
- [35] N. E. Richards and W. D. Smart, “How Should the Law Think About Robots?” *University of Washington Law Review*, vol. 86, pp. 119–160, 2013.
- [36] T. Z. Zarsky, “Algorithmic Decision-Making and the Right to Explanation,” *Berkeley Technology Law Journal*, vol. 34, pp. 189–235, 2019.
- [37] F. Schauer, *Profiles, Probabilities, and Stereotypes*, Harvard University Press, 2003.
- [38] M. Veale and I. Brass, “Administration by Algorithm? Public Management Meets Public Sector Machine Learning,” *Public Management Review*, vol. 21, no. 12, pp. 2050–2072, 2019.
- [39] B. Goodman and S. Flaxman, “European Union Regulations on Algorithmic Decision-Making and a ‘Right to Explanation’,” *AI Magazine*, vol. 38, pp. 50–57, 2017.
- [40] D. Freeman, “Constitutional Fairness and AI Bias,” *Harvard Civil Rights–Civil Liberties Law Review*, vol. 57, pp. 261–301, 2022.
- [41] M. A. Lemley and B. Casey, “Remedies for Robots,” *University of Chicago Law Review*, vol. 86, pp. 1311–1396, 2019.
- [42] E. Engstrom and D. Ho, “Algorithmic Accountability in the Administrative State,” *Yale Journal on Regulation*, vol. 37, pp. 800–848, 2020.
- [43] P. Hacker, “Teaching Fairness to Artificial Intelligence: Existing and Novel Strategies Against Algorithmic Discrimination,” *Berkeley Technology Law Journal*,

vol. 35, pp. 113–164, 2020.

[44] R. E. Boyd, “Equal Protection, AI, and the Color of Data,” *Columbia Journal of Law & Social Problems*, vol. 55, pp. 101–146, 2022.

[45] J. Obermeyer et al., “Dissecting Racial Bias in an Algorithm Used to Manage Health Populations,” *Science*, vol. 366, no. 6464, pp. 447–453, 2019.

[46] J. A. Kroll et al., “Accountable Algorithms,” *University of Pennsylvania Law Review*, vol. 165, pp. 633–705, 2017.

[47] K. E. Davis, “Constitutionalizing Artificial Intelligence: Judicial Review in an Algorithmic Age,” *American University Law Review*, vol. 70, pp. 1041–1097, 2021.

[48] S. D. Warren and L. D. Brandeis, “The Right to Privacy,” *Harvard Law Review*, vol. 4, no. 5, pp. 193–220, 1890.

[49] O. Keyes, “The Misgendering Machines: Trans/HCI Implications of Automatic Gender Recognition,” *Proceedings of the ACM Human Factors in Computing Systems (CHI)*, 2018.

[50] A. Ezrach, “Algorithmic Equality and State Action,” *Harvard Journal of Law & Technology*, vol. 35, pp. 1–38, 2022.

[51] M. Sandel, *Justice: What’s the Right Thing to Do?*, Farrar, Straus and Giroux, 2010.

[52] R. Dworkin, *Taking Rights Seriously*, Harvard University Press, 1977.

[53] L. Fuller, *The Morality of Law*, Yale University Press, 1969.

[54] A. D. Selbst, “Disparate Impact in Big Data Policing,” *Georgia Law Review*, vol. 52, pp. 109–195, 2017.

[55] K. Mathias, “Beyond Formal Neutrality: Equal Protection and Algorithmic Discrimination,” *UCLA Law Review Discourse*, vol. 68, pp. 201–234, 2021.