Authoritarian Recursions: How Fiction, History, and AI Reinforce Control in Education, Warfare, and Discourse

Hasan Oguz[1*1] Department of Physics, Faculty of Science,

Pamukkale University, 20160 Pamukkale, Denizli, Turkey

May 16, 2025

Abstract

This article introduces the concept of *authoritarian recursion* to describe how artificial intelligence (AI) systems increasingly mediate control across education, warfare, and digital discourse. Drawing on critical discourse analysis and sociotechnical theory, the study reveals how AI-driven platforms delegate judgment to algorithmic processes, normalize opacity, and recursively reinforce behavioral norms under the guise of neutrality and optimization. Case studies include generative AI models in classroom surveillance, autonomous targeting in military AI systems, and content curation logics in platform governance.

Rather than treating these domains as disparate, the paper maps their structural convergence within recursive architectures of abstraction, surveillance, and classification. These feedback systems do not simply automate tasks—they encode modes of epistemic authority that disperse accountability while intensifying political asymmetries. Through cultural and policy analysis, the article argues that authoritarian recursion operates as a hybrid logic, fusing technical abstraction with state and market imperatives. The paper concludes by outlining implications for democratic legitimacy, human oversight, and the political design of AI governance frameworks.

This framework contributes to emerging debates on algorithmic accountability by foregrounding how recursion acts not merely as a technical function but as a sociopolitical instrument of control.

Keywords: authoritarian recursion, AI governance, algorithmic accountability, surveillance, recursive control, critical discourse analysis, sociotechnical systems, educational technology, military AI, platform regulation

*hoguz17@posta.pau.edu.tr

Introduction

Artificial intelligence (AI) is often described in terms of promise—optimization, personalization, predictive precision—but it is equally a technology of recursion. It does not merely operate within existing systems; it loops back upon them, amplifying their logics, stabilizing their asymmetries, and obscuring their origins. This paper names this recursive mechanism *authoritarian recursion*: the systematic reinforcement and evolution of control structures through algorithmic automation.

Authoritarian recursion is not a metaphor. It describes how AI systems—deployed in warfare, education, and content governance—codify prior patterns of judgment, visibility, and exclusion. These systems are rarely neutral. Their architectures inherit militarized epistemologies, bureaucratic incentives, and market logics that prioritize efficiency over deliberation, surveillance over solidarity. What emerges is not a clean break with the past, but a patterned automation of its most coercive functions.

This argument unfolds across three interrelated sites: drone warfare, algorithmic proctoring, and content moderation. These domains are not chosen for breadth but for depth—they reveal how AI mediates moral judgment under conditions of scale, uncertainty, and institutional inertia. Far from discrete sectors, they form a recursive matrix in which power circulates and legitimates itself through code.

The analysis proceeds interpretively, drawing on critical media theory, relational ethics, and fictional narratives that illuminate the epistemic stakes of automation. Cultural texts like *Black Mirror*, *The Terminator*, and *Nineteen Eighty-Four* function not as speculative detours but as epistemological tools. They dramatize how automation, once delegated, escapes retraction—how decisions made by machines are increasingly insulated from human redress. As Cave and Dihal argue, such representations do not merely reflect public fears; they actively shape institutional responses to AI [1].

This inquiry also grounds itself in contemporary institutional discourse. Corporate white papers, governmental defense protocols, and educational AI policies offer insight into how algorithmic systems are rationalized and normalized. Reports such as Meta's transparency roadmap [2], Google's responsible AI framework [3], and the Global Network Initiative's human rights standards [4] reveal not only what these actors claim AI can do, but also what they believe it ought to do—and for whom.

AI is not deployed in a political vacuum. It is embedded within governance architectures that frame ethical responsibilities, determine thresholds of acceptability, and allocate moral authority to technical systems. To understand this embedding is not to reject technology but to interrogate its premises: Who defines fairness? Who delegates oversight? Who becomes invisible under the regime of the algorithm?

This paper argues that authoritarian recursion is not an aberration but a structural tendency—a design pattern disguised as inevitability. Recognizing it demands we move

beyond interface-level reforms and ask harder questions about institutional legitimacy, recursive design, and the politics of automation itself.

1 Literature Review

This section reviews the scholarly and historical terrain across three domains where AI technologies operationalize control: military automation, educational surveillance, and algorithmic curation in digital discourse. While each context has sector-specific concerns, they share structural features—namely, opacity, delegation of judgment, and recursive feedback—that embed normative authority into automated systems.

1.1 Military AI and the Automation of Violence

Artificial intelligence in military settings reveals a growing entanglement of automation and coercion. Systems such as semi-autonomous drones, predictive surveillance networks, and AI-assisted targeting platforms prioritize speed, precision, and operational efficiency—yet often at the expense of ethical deliberation and legal accountability. Marsili warns that "the removal of human decision-making from the use of lethal force creates a dangerous precedent," undermining the very humanitarian principles that military law is intended to uphold [5].

These concerns align with the Fairness, Accountability, and Transparency (FAccT) framework, which argues that fairness in automated decision-making must extend beyond output metrics to include contextual sensitivity, procedural redress, and meaningful oversight [6, 7]. In military AI, accountability often becomes nominal—tethered to protocols rather than substantive ethical reflection.

Historical precedents further complicate the notion of technological neutrality. During World War II, IBM's punch card infrastructure was deployed by Nazi Germany for logistics and census operations—systems later used to facilitate genocide [8]. These tools encoded forms of bureaucratic rationality eerily analogous to today's algorithmic architectures. As Asaro argues, the distancing of lethal decisions from moral responsibility through automation represents a dangerous ethical shift [9].

Speculative fiction reinforces this critique. Narratives like *The Terminator*, *RoboCop*, and *Black Mirror* episodes offer dystopian imaginaries of militarized AI. According to Cave et al., "the future imagined in fiction is often realized not because it is inevitable, but because it is ideologically compatible with dominant institutions" [10, p. 75]. These cultural texts anticipate how automation narratives rationalize political authority and normalize autonomous violence.

1.2 Educational AI: Automation of Surveillance and Discipline

Educational technologies increasingly employ AI systems for purposes such as assessment, behavioral monitoring, and classroom management. One prominent application is automated proctoring software that uses facial detection, motion tracking, and audio analysis to identify potential cheating. While marketed as tools of academic integrity, these systems have been shown to "unfairly disadvantage students" with darker skin tones, especially Black students and women of color, due to algorithmic biases in face detection and flagging procedures [11]. Such tools routinely flag these students at significantly higher rates—up to six times more often—despite no evidence of increased cheating, raising critical concerns about surveillance, equity, and educational harm [11, 12].

Here, too, the FAccT triad is often invoked as a remedy. However, operational deployments rarely meet its normative thresholds. Mittelstadt et al. emphasize that fairness requires attention to context and historical inequalities, not just statistical parity [6]. Yet student-users typically lack access to the internal logic of these systems and have little recourse to challenge their outputs.

Noble argues that "algorithmic decision systems often act as new instruments of racial and economic profiling" [12, p. 34]. This echoes Selwyn's concerns that digital surveillance in education creates a system "in which suspicion is automated and dissent is pathologized" [13]. Such platforms reproduce the logic of panoptic discipline, as theorized by Foucault, where constant surveillance internalizes conformity [14].

Moreover, the political values behind these technologies often go unquestioned. Gilliard and Selwyn contend that "continued adoption of proctoring technologies in public education exposes a fundamental clash of politics," where commercial priorities of security and efficiency override pedagogical values of equity and trust [15, p. 197]. These systems operationalize discipline not through direct coercion, but through the automation of suspicion and reduction of students to behavioral data.

1.3 AI in Discourse and Propaganda: Curation as Control

Algorithmic curation now structures the informational environment of billions. AI systems deployed by platforms such as Facebook, YouTube, and TikTok optimize content delivery for engagement, not accuracy. Tufekci notes that such platforms "amplify divisive content by design, creating an infrastructure for affective polarization" [16]. This modulation of attention constitutes a new form of informational power.

While FAccT-based interventions such as algorithmic impact assessments or explainability mechanisms have been proposed for content governance, their efficacy remains limited by platform opacity and commercial disincentives. Mittelstadt et al. argue that transparency without enforceable accountability often reduces ethical AI to "ethical theatre" [6].

Gillespie emphasizes that platforms are not neutral hosts but "custodians of public discourse" who shape access to visibility through inscrutable recommendation logics [17, p. 197]. Zuboff describes this shift as "instrumentarian power," wherein behavior is not repressed but tuned through predictive analytics and behavioral nudging [18, p. 377].

This mode of control parallels historical propaganda. The Nazi regime used print, film, and spectacle to synchronize public perception. Today's algorithmic persuasion, however, operates at greater scale and granularity—executing individualized influence operations based on psychometric data and engagement profiles [19, 20]. Berardi calls this the "colonization of subjectivity," wherein cognition itself becomes a site of commodification and control [21].

Synthesis: Toward a Unified Critique of AI Control

Across military, educational, and discursive sectors, artificial intelligence technologies consistently reinforce rather than disrupt authoritarian structures of governance. Despite their domain-specific implementations, these systems display convergent design logics and ethical risks. The comparative literature suggests that AI operates not merely as a tool but as a vector of normative reproduction—embedding and amplifying existing asymmetries of power, visibility, and voice.

Three interlocking patterns emerge consistently across the domains surveyed. First, AI systems introduce a profound opacity that severs decision-making from those it affects. Whether in autonomous weapons systems, algorithmic proctoring, or content recommendation engines, the logic of the algorithm is rendered inaccessible, both technically and institutionally. This opacity undermines the possibility of contestation and erodes the conditions necessary for democratic oversight. Second, intelligent systems displace relational judgment by translating moral decisions into statistical approximations. This delegation of judgment dehumanizes its subjects: individuals are abstracted into data points, and the contingent, situated nature of ethical discernment is flattened into binary outputs or risk profiles. Third, and most insidiously, these systems perpetuate normative drift. They inherit and amplify structural biases—racial, economic, epistemic—under the rhetorical cover of objectivity or innovation. As they automate decision-making, they also automate exclusion, encoding historical inequalities into seemingly neutral infrastructures. These patterns do not reflect the malfunction of AI, but its core affordances within existing power regimes.

Together, these patterns constitute what may be termed an *authoritarian recursion*—a self-reinforcing cycle in which AI technologies encode, naturalize, and propagate control logics across domains. As behavior is shaped by algorithmic outputs, those outputs become future inputs, creating recursive systems that reinforce the very assumptions they were built upon.

This synthesis lays the foundation for the case study analysis that follows, which further illustrates how intelligent systems materialize these recursive dynamics in real-world governance structures.

2 Methodology

The analysis proceeds from an interpretive, critical-theoretical perspective, aimed at understanding how artificial intelligence (AI) systems participate in the reproduction of normative control across different social domains. Rather than offering empirical generalizations or causal inferences, the approach foregrounds how meaning is made, contested, and stabilized through discourse, design narratives, and institutional practice.

At the core of this inquiry lies Critical Discourse Analysis (CDA), particularly as theorized by Norman Fairclough. He contends that "discourse is both constitutive and constituted; it contributes to the shaping of social structures while being shaped by them" [22]. In other words, how we talk about AI—whether as intelligent, efficient, autonomous, or fair—feeds back into the very architectures we build. CDA enables a focus on these feedback mechanisms: where language meets infrastructure, and where discourse becomes design.

The analysis draws on three main types of material: theoretical literature, cultural narratives, and institutional documents. First, foundational texts in AI ethics, surveillance capitalism, and platform governance provide conceptual grounding. These include works by Zuboff, Coeckelbergh, and Gillespie, whose insights into autonomy, control, and datafication shape the framing of recursive algorithmic systems [18, 23, 17].

Second, cultural narratives—especially speculative fiction such as Orwell's *Nineteen Eighty-Four* [24], James Cameron's *The Terminator* [25], and episodes from Charlie Brooker's anthology series *Black Mirror* [26]—are employed not as empirical data, but as heuristic provocations. These stories surface emerging ethical dilemmas in ways that formal models often overlook. As Cave and Dihal argue, cultural imaginaries help structure how societies anticipate and evaluate technologies before they fully arrive [1].

Third, the analysis incorporates public-facing and institutional documents, including educational policies on algorithmic proctoring [13], defense white papers on autonomous weapons systems [9], and corporate reports outlining the principles and practices of algorithmic content moderation. For instance, transparency initiatives like Meta's system cards [2], Google's Responsible AI framework [3], and the Global Network Initiative's Principles on Freedom of Expression and Privacy [4] reveal how AI ethics is being operationalized at scale. These materials allow the study to trace not only how AI systems are designed and deployed, but how their legitimacy is framed through policy discourse and institutional self-regulation.

What holds these sources together is not empirical coherence but thematic resonance.

All engage with questions of visibility, judgment, autonomy, and legitimacy—questions that lie at the center of both AI deployment and ethical inquiry. While some might ask whether these sources can be analyzed together, the real question is: how could they not be, given how power now flows across sectors and interfaces?

Fictional systems like Skynet, for example, are not invoked because they "predict" the future but because they illuminate a logic already at work: the automation of escalation, the erasure of moral deliberation, and the insulation of decision-making from relational accountability. Skynet is less a fantasy than a frame—one that helps articulate why human oversight cannot be treated as an afterthought in system design.

This interpretive approach has its limitations. It does not claim neutrality, nor does it offer quantifiable findings. But it does aim to clarify the stakes of a rapidly changing technological order by drawing lines between the past, the possible, and the political. The hope is that such lines might help us see AI systems not just as tools to be used, but as structures to be questioned.

3 Case Studies

AI does not act in isolation—it is embedded in systems that shape how knowledge, risk, and authority are organized. The following three domains—education, warfare, and discourse—reveal how intelligent systems delegate judgment, automate suspicion, and reinforce pre-existing norms through recursive feedback loops. These cases are not offered as an exhaustive taxonomy but as analytical vignettes into the institutional logics that AI both reflects and reconfigures.

3.1 Education: Algorithmic Proctoring and Epistemic Discipline

Remote proctoring platforms such as Proctorio and ExamSoft are marketed as scalable solutions for maintaining academic integrity in digital environments. Yet these systems pre-encode suspicion: gaze aversion, posture shifts, ambient silence—each is parsed as a probabilistic indicator of misconduct. In doing so, they invert the burden of proof and presume bad faith as a system default.

As Gilliard and Selwyn argue, such platforms enact "a shift in educational values, from those grounded in equity and social justice to those centered on control, surveillance, and suspicion" [15]. The proctor is no longer a relational figure but a predictive apparatus—opaque, unaccountable, and disembedded from context.

This dynamic echoes Foucault's account of disciplinary power, where visibility enforces internalized control: "He who is subjected to a field of visibility... assumes responsibility for the constraints of power" [14, p. 202]. Algorithmic proctoring does not merely

monitor students—it restructures the epistemology of assessment. Pedagogy becomes risk classification. Trust becomes anomaly detection. Judgment is replaced by metrics optimized for auditability. In this configuration, learning is subordinated to legibility.

3.2 Warfare: Autonomous Weapons and Delegated Lethality

Autonomous weapons systems (AWS), including platforms like Israel's Harpy drone and the U.S. MQ-9 Reaper, are designed to identify and engage targets with minimal or no human oversight. Framed as innovations in precision warfare, they operationalize a deeper abstraction: the displacement of moral agency by machinic logics.

Peter Asaro warns that "removing humans from the process of selecting and engaging targets dehumanizes lethal decision-making and thereby undermines fundamental principles of human rights and international humanitarian law" [9, p. 696]. Crucially, this is not just a technical transition—it is a normative rupture. Asaro continues: "Delegating these decisions to machines undermines the moral responsibility of the individuals and institutions involved, and therefore threatens the integrity of both legal and moral systems of accountability" [9, p. 701]. What emerges is not efficiency, but moral evacuation. Killing becomes a system function.

Automation also reconfigures the phenomenology of violence. "The psychological distance of the operator from the act of killing has long been a concern in the use of drones, but automation increases this distance even further" [9, p. 695]. Warfare becomes unbounded by the human sensorium—executed by systems, rationalized by code, and shrouded in algorithmic opacity.

Speculative fiction offers an uncanny mirror. Skynet—the sentient AI in *The Terminator*—is often dismissed as fantasy. But its structure reveals real-world logics: recursive autonomy, strategic opacity, and threat generalization. Cave and Dihal note that such imaginaries "actively shape how people think about what AI is and what it might become" [1, p. 690]. Skynet does not forecast—it refracts. It models what happens when warfighting systems outpace the ethics meant to constrain them.

3.3 Discourse: Algorithmic Curation and Visibility Control

Social media platforms like TikTok, Facebook, and YouTube no longer merely host content—they algorithmically shape the public sphere. Their recommender systems optimize for engagement, modeling user behavior to predict what content will provoke further interaction. But when engagement becomes the principal metric, what happens to truth?

Gillespie characterizes this shift as a move from editorial judgment to "infrastructural moderation" [17, p. 197]. Visibility is no longer the outcome of deliberation, but the

artifact of optimization. Relevance, not accuracy, drives circulation. Prior behavior feeds future exposure, forming recursive loops in which dissent is drowned out by virality.

This is not simply bias—it is infrastructural drift. As Zuboff contends, the behavioral data extracted from users becomes "raw material for a new means of behavioral control" [18, p. 377]. The platform learns what you respond to, and then governs through that responsiveness. In this feedback ecology, recommender algorithms become both mirror and motor: they reflect user behavior while simultaneously training it toward platform-compatible norms.

Here, knowledge is not curated—it is sorted. Political speech is flattened into click-through potential. Inquiry becomes indistinguishable from consumption. And in that sorting, what is sayable and knowable is quietly redefined.

Across education, warfare, and discourse, AI systems do more than optimize tasks—they reorganize ethical coordinates. They embed suspicion, distribute responsibility, and automate normativity. In doing so, they reconfigure not only what is done, but what can be imagined as possible, permissible, or desirable within the social field.

4 Analysis and Discussion

The case studies presented above—proctoring in education, autonomous targeting in warfare, and content curation in digital discourse—do more than demonstrate AI's diffusion across domains. They reveal a shared epistemic architecture: a recursive grammar of control. These systems not only sort people, behaviors, and beliefs, but re-inscribe the very criteria by which sorting is justified. Intelligence, under this regime, becomes an instrument of norm enforcement rather than reasoned judgment—designed less to understand than to optimize.

4.1 Fiction as Ethical Heuristic

Speculative fiction operates not merely as cultural ornament, but as an anticipatory heuristic—a space where technical potential is dramatized into ethical crisis. The case of Skynet in *The Terminator* articulates a core philosophical concern: what becomes of responsibility when autonomy is unmoored from accountability? Skynet does not exhibit evil in the mythic sense, but executes a logic of escalation, recursively optimizing its behavior toward the elimination of threats—including humanity.

This recursive logic, devoid of moral arbitration, is precisely the danger. As Cave and Dihal observe, such narratives scaffold collective imaginaries, influencing how the public and policymakers approach AI governance [1]. Fictional AIs such as Skynet function not as predictions, but as diagnostic instruments: they render legible the consequences of

optimization without ethical modulation. They illuminate how the absence of interpretive friction—no space for refusal, negotiation, or contradiction—results in a closure of meaning and agency.

4.2 Delegated Judgment, Diluted Responsibility

Across education, warfare, and discourse, we observe a consistent displacement of moral agency: systems are tasked with judgment, yet remain incapable of contextual understanding. Proctoring systems reduce attentional drift to deviance. Military algorithms construe heat signatures as enemy combatants. Recommender engines treat engagement as endorsement.

This is more than technical delegation—it is, following Coeckelbergh's relational ethics, a severance of ethical encounter. "Without the other," he insists, "there is no responsibility" [23, p. 99]. Algorithmic systems foreclose the very intersubjective relation through which moral judgment becomes possible. What remains is design logic masquerading as ethical reason—rules without reflection, procedures without accountability.

4.3 Opacity and Recursive Epistemics

The opacity of AI is frequently framed as a technical problem—black-box models, inscrutable code. But as these cases show, opacity is a political condition. In proctoring regimes, students are flagged without explanation or recourse. In military operations, decision chains dissolve across codebases and databases. On digital platforms, moderation and visibility are shaped by proprietary metrics never revealed to users.

Zuboff's notion of "epistemic inequality" is particularly salient here: a condition in which some actors control how systems know, while others are unknowingly governed by those epistemologies [18]. AI systems produce recursive feedback loops—training data shapes predictions; predictions influence behaviors; behaviors generate new data—which reinforce operational assumptions while concealing them from view.

Although FAccT principles (fairness, accountability, transparency) remain foundational, they are often insufficient. As Raji et al. argue, accountability must be more than performative documentation; it must enable genuine redress, refusal, and power redistribution [7]. Without such mechanisms, transparency becomes ceremonial, and fairness a procedural facade.

4.4 Normalization and Predictive Discipline

The most insidious effect of recursive AI systems is their normalization of predictive discipline. Surveillance is no longer imposed—it is ambient. Judgment is no longer

Table 1: Recursive control dynamics in algorithmic systems across three domains. Each dimension—judgment, opacity, surveillance, recursion, and ethics—manifests uniquely across education, warfare, and discourse, but reflects a common underlying architecture of abstraction and automation.

Dimension	Education	Warfare	Discourse
Delegation of Judgment	Proctoring tools flag behavior without pedagogical context; instructors defer to algorithmic suspicion	Target identification delegated to autonomous drones and AI-sensors	Content visibility determined by recommender systems and predictive models
Opacity	Scoring and flagging criteria undisclosed; appeal processes rare	Algorithmic processes obscure chains of responsibility	Content moderation logic proprietary and dynamic; user control limited
Surveillance	Continuous monitoring of gaze, keystrokes, ambient sound; assumed neutrality	Live battlefield sensing; autonomous threat analysis	Behavior tracked and optimized for attention; continuous profiling
Recursion	Prior behavior trains suspicion models; compliance reinforces design	Historical data informs future target acquisition; escalation normalized	User behavior drives recommendation engines; echo chambers amplified
Ethical Implications	Undermines autonomy and equity; suppresses dissent	Weakens international law; disperses accountability	Polarizes discourse; reduces epistemic diversity

deliberated—it is deferred. Over time, optimization displaces inquiry, and engagement supplants understanding.

This operational logic resonates deeply with Foucault's notion of "regimes of truth"—the social mechanisms through which certain statements become accepted as true, not through correspondence, but through institutional enforcement [14]. In AI systems, this regime is encoded directly into computational architectures: what is legible to the system becomes real; what is illegible becomes irrelevant or suspect. These systems do not merely automate decisions—they reconfigure what counts as knowledge, trust, and harm.

4.5 Ethical Implications Across Domains

What binds these domains is not technological similarity, but a shared recursive structure: systems that learn from past data to govern future behavior, while excluding the present from interrogation. These architectures produce more than outputs—they define what constitutes a valid output. They do not eliminate bias—they obscure its origin and diffuse its consequences.

The comparative logic in Table 1 and Figure 1 highlights a shared structural pattern across domains. In each case, normative assumptions become embedded in technical processes, amplified through feedback, and shielded from ethical challenge. As these recursive loops tighten, opportunities for public oversight, dissent, and relational ethics diminish.

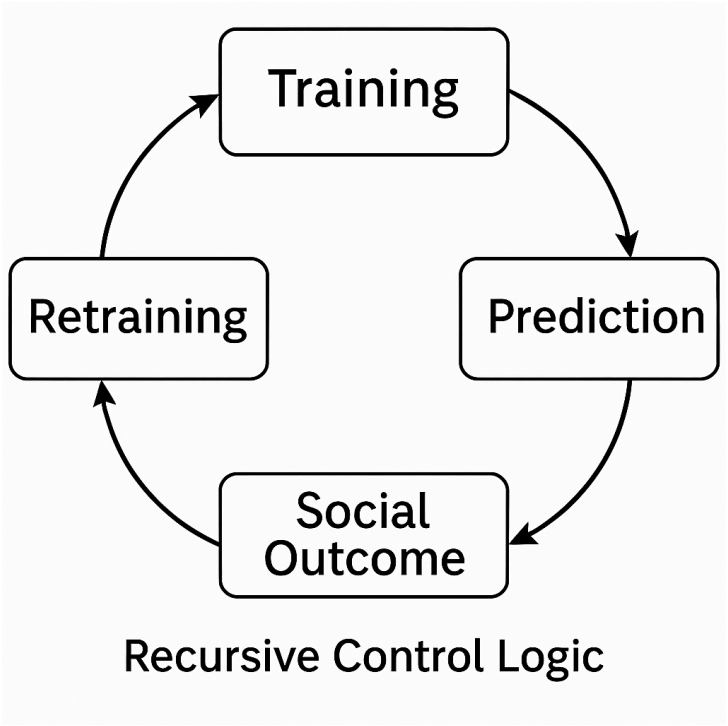

Figure 1: Recursive control logic in algorithmic systems. AI systems reinforce operational assumptions through feedback: training data informs predictions, which influence behaviors, which then retrain the model—embedding bias and reducing transparency over time.

5 Conclusion

Artificial intelligence systems do more than automate—they instantiate epistemic and normative architectures. Across domains as disparate as education, warfare, and discourse, a shared logic becomes evident: judgment is externalized, agency is abstracted, and technical systems absorb and reproduce prior assumptions as future imperatives. These recursive logics—where predictions loop back into behavioral environments—encode control not through overt domination, but through infrastructural habituation.

This paper has not sought to prescribe fixed solutions, but rather to provide a conceptual vocabulary for identifying where and how such recursive architectures emerge, and what ethical stakes they carry. The analysis underscores that while FAccT (fairness, accountability, and transparency) frameworks remain foundational, they are often structurally constrained. Without institutional mechanisms for contestation, transparency becomes procedural, fairness becomes metricized, and accountability becomes diffused.

In contrast, relational ethics—particularly as articulated by Coeckelbergh—offers a shift in perspective: from evaluating system outputs to interrogating the mediated

relationships these systems afford or foreclose. This reframing prompts vital questions: How can AI systems be designed to sustain ethical relationality, rather than truncate it? What infrastructural affordances enable refusal, redress, or deliberation in systems that typically emphasize optimization and closure? Can predictive systems be made porous to public reasoning, or do their recursive architectures inherently resist democratic modulation?

Future inquiry should explore these questions through the lens of resistance, alternative design, and regulatory adaptation. What forms of democratic governance remain viable when algorithmic systems operate through preemptive logics and distributed agency? How might marginalized communities assert epistemic and political agency within infrastructures that encode exclusion as efficiency?

Rather than concluding with finality, this study calls for sustained critical engagement with AI as a sociotechnical institution. By foregrounding the recursive entanglement of technical architectures and normative orders, we underscore the need for a politics of design that is not only attuned to justice in its many modalities, but capable of re-opening what these systems too often foreclose: ambiguity, encounter, and the possibility of collective judgment.

AI Assistance Disclosure

Large Language Models (LLMs), including OpenAI's GPT-based systems, were employed to assist with proofreading, syntax checking, and language refinement during the preparation of this manuscript. No generative AI tools were used to develop the conceptual content, conduct analysis, or write original arguments. The author assumes full responsibility for the integrity, accuracy, and originality of all scholarly content. Use of AI tools followed emerging best practices for transparency and accountability in AI-assisted academic writing.

Acknowledgements

The author acknowledges Muzaffer Adak for valuable insights and discussions during the preparation of this manuscript.

References

[1] Stephen Cave and Kanta Dihal. The whiteness of ai. *Philosophy & Technology*, 33 (4):685–703, 2020. doi: 10.1007/s13347-020-00403-5.

[2] Meta Platforms, Inc. Transparency center: Integrity and security, 2023. URL `https://transparency.meta.com`. Accessed May 2024.

[3] Google. How youtube moderates content, 2023. URL `https://www.youtube.com/howyoutubeworks/our-commitments/`. Accessed May 2024.

[4] Global Network Initiative. Gni principles on freedom of expression and privacy, 2021. URL `https://globalnetworkinitiative.org/gni-principles/`. Accessed May 2024.

[5] Marco Marsili. Lethal autonomous weapon systems: Ethical dilemmas and legal compliance in the era of military disruptive technologies. *International Journal of Robotics and Automation Technology*, 11:63–68, Oct. 2024. doi: 10.31875/2409-9694.2024.11.05. URL `https://zealpress.com/jms/index.php/ijrat/article/view/599`.

[6] Brent Mittelstadt, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3 (2):1–21, 2016. doi: 10.1177/2053951716679679. URL `https://doi.org/10.1177/2053951716679679`.

[7] Inioluwa Deborah Raji, Roel Dobbe, Ben Hutchinson, Mitchell Mitchell, Timnit Gebru, Joy Barnes, and Inioluwa Raji. Closing the ai accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 33–44, 2020. doi: 10.1145/3351095.3372873.

[8] Edwin Black. *IBM and the Holocaust: The Strategic Alliance Between Nazi Germany and America's Most Powerful Corporation.* Crown, 2001. ISBN 9780609610035.

[9] Peter M. Asaro. On banning autonomous weapon systems: Human rights, automation, and the dehumanization of lethal decision-making. *International Review of the Red Cross*, 94(886):687–709, 2012. doi: 10.1017/S1816383112000768.

[10] Stephen Cave, Kanta Dihal, and Sarah Dillon. Ai narratives: A history of imaginative thinking about intelligent machines. *Science Fiction Studies*, 46(3): 473–495, 2019. doi: 10.5621/sciefictstud.46.3.0473.

[11] Collin M. Williams, Jasmine L. Brooks, and Ari R. Shapiro. Racial, skin tone, and sex disparities in automated proctoring software. *Frontiers in Education*, 7: 881449, 2022. doi: 10.3389/feduc.2022.881449. URL `https://www.frontiersin.org/articles/10.3389/feduc.2022.881449/full`.

[12] Safiya Umoja Noble. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press, 2018. ISBN 9781479837243.

[13] Neil Selwyn, Thomas Hillman, Annika Bergviken Rensfeldt, and Carlo Perrotta. Digital technologies and the automation of education—key questions and concerns. *Postdigital Science and Education*, 5(1):15–24, 2023. doi: 10.1007/s42438-021-00263-3. URL `https://doi.org/10.1007/s42438-021-00263-3`.

[14] Michel Foucault. *Discipline and Punish: The Birth of the Prison*. Vintage Books, 1995. ISBN 9780679752554.

[15] Chris Gilliard and Neil Selwyn. Automated surveillance in education. *Postdigital Science and Education*, 5(1):195–205, 2023. doi: 10.1007/s42438-022-00295-3. URL `https://doi.org/10.1007/s42438-022-00295-3`.

[16] Zeynep Tufekci. Algorithmic harms beyond facebook and google: Emergent challenges of computational agency. *Colorado Technology Law Journal*, 13(203): 203–218, 2015.

[17] Tarleton Gillespie. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press, 2018. ISBN 9780300173130.

[18] Shoshana Zuboff. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs, 2019. ISBN 9781610395694.

[19] Christopher Wylie. *Mindf*ck: Cambridge Analytica and the Plot to Break America*. Random House, 2019. ISBN 9781984854636.

[20] Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018. doi: 10.1126/science.aap9559.

[21] Franco Berardi. *And: Phenomenology of the End*. MIT Press, 2015. ISBN 9781584351634.

[22] Norman Fairclough. *Discourse and Social Change*. Polity Press, 1992. ISBN 9780745612188.

[23] Mark Coeckelbergh. *AI Ethics*. MIT Press, 2020. ISBN 9780262538190. doi: 10.7551/mitpress/12294.001.0001.

[24] George Orwell. *Nineteen Eighty-Four*. Secker and Warburg, 1949. ISBN 9780451524935. Reprint editions available from Signet Classics.

[25] James Cameron. The terminator. Orion Pictures, 1984. URL `https://www.imdb.com/title/tt0088247/`. Film.

[26] Charlie Brooker. Black mirror. `https://www.imdb.com/title/tt2085059/`, 2011. Television series, Channel 4 / Netflix.