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

Hasan Oguz*

Department of Physics, Faculty of Science,

Pamukkale University, 20160 Pamukkale, Denizli, Turkey

June 8, 2025

Abstract

This article develops the concept of *authoritarian recursion* to theorize how artificial intelligence (AI) systems consolidate institutional control across education, military operations, and digital discourse. Rather than treating these domains in isolation, it identifies a shared recursive architecture in which algorithmic systems mediate judgment, obscure accountability, and reshape the conditions of moral and epistemic agency.

Grounded in critical discourse analysis and sociotechnical ethics, the paper synthesizes historical precedent, cultural narrative, and contemporary deployment to examine how intelligent systems normalize hierarchy under the guise of efficiency and neutrality. Case studies include automated proctoring in education, autonomous targeting in warfare, and algorithmic curation on social platforms. Cultural imaginaries such as Orwell's *Nineteen Eighty-Four*, *The Terminator*'s Skynet, and *Black Mirror* are treated as heuristic devices that illuminate public anxieties and design assumptions embedded in technological systems.

The analysis integrates frameworks from the Fairness, Accountability, and Transparency (FAccT) paradigm, relational ethics, and data justice theory to explore the normative implications of predictive infrastructures. It argues that recursive control operates through moral outsourcing, behavioral normalization, and epistemic closure. By reframing AI not as a neutral tool but as a communicative and institutional infrastructure, the article highlights the need for ethical orientations that prioritize democratic refusal, epistemic plurality, and responsible design in the governance of intelligent systems.

*hoguz17@posta.pau.edu.tr

Keywords: authoritarian recursion, algorithmic governance, recursive control, predictive infrastructures, sociotechnical ethics, platform power, educational surveillance, autonomous weapons, content moderation, critical discourse analysis, data justice, relational responsibility, infrastructural opacity, epistemic closure, democratic refusal

1 Introduction

Artificial intelligence (AI) is no longer merely a domain of technical optimization; it increasingly operates as an infrastructure of communication and governance. From remote proctoring platforms in education to autonomous weapons and content moderation systems, AI technologies are central to how institutions classify, interpret, and act on information. These systems do not simply automate decisions—they shape epistemic boundaries, delegate judgment, and redistribute communicative agency. In doing so, they embed themselves within broader sociotechnical arrangements of power and legitimacy.

This paper examines AI as a recursive mode of governance: a system that learns from, intervenes in, and ultimately reshapes the very behaviors it observes. These feedback architectures operate not only through prediction and abstraction, but through communicative invisibility. As platforms refine recommendations, military systems automate targeting, and educational software models behavioral suspicion, recursive AI infrastructures begin to enact what Couldry and Mejias term "data colonialism"—the capture and reconfiguration of human life through extraction and abstraction [1]. They reflect Gillespie's notion of "infrastructural power," where visibility and participation are mediated through algorithmic protocols [2], and contribute to what van Dijck et al. describe as "platform governance"—a mode of rule exercised through technical standards and informational architectures [3].

To theorize these dynamics, the article introduces the concept of *authoritarian recursion*. This term describes the self-reinforcing loops through which intelligent systems encode, legitimize, and propagate control logics—frequently under the rhetorical cover of personalization, neutrality, or operational efficiency. These recursive infrastructures obscure responsibility, foreclose contestation, and deepen asymmetries between users, institutions, and machinic systems.

The argument proceeds through three case domains—education, warfare, and digital discourse—where recursive AI systems materialize different forms of delegated authority and normative closure. Drawing on critical discourse analysis, the study treats AI not simply as a tool, but as a communicative actor that shapes who is visible, what is knowable, and which actions are thinkable. Cultural imaginaries such as *Black Mirror* and *The Terminator*'s Skynet are incorporated not as evidence, but as critical heuristics that reflect and amplify public anxieties around machinic autonomy and recursive control.

In dialogue with critical media studies, surveillance scholarship, and sociotechnical critique [4, 5, 6], the article advances the study of algorithmic governance by foregrounding recursion as both a technical logic and an ideological formation. Rather than proposing universal solutions or narrow design principles, the paper advocates for renewed scrutiny of how AI systems enact communicative authority—structuring infrastructures of attention, accountability, and legitimacy across institutional fields.

2 Literature Review

This section synthesizes scholarship on how AI technologies function as instruments of sociotechnical control in three domains: military automation, educational surveillance, and digital discourse. These sectors are often treated as distinct, yet their AI applications share design logics—opacity, delegated authority, and recursive feedback—that normalize institutional power under the guise of optimization. This literature provides the groundwork for theorizing authoritarian recursion as a mode of governance embedded in platforms and predictive systems. Drawing from ethical frameworks, historical precedents, and critical media theory, the review integrates both empirical and conceptual contributions to illustrate how automated systems entrench and legitimize normative authority across contexts.

2.1 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 [7].

These concerns align with the Fairness, Accountability, and Transparency (FAccT) framework, which holds that fairness in automated decision-making must extend beyond output metrics to include contextual sensitivity, procedural redress, and meaningful oversight [8, 9]. 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 [10]. These tools exemplified a bureaucratic rationality disturbingly resonant with today's algorithmic architectures. As Asaro argues, the distancing of lethal decisions from moral responsibility

through automation represents a dangerous ethical shift [11].

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" [12, p. 75]. These cultural texts anticipate how automation narratives rationalize political authority and normalize autonomous violence.

2.2 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 [13]. 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 [13, 14].

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 [8]. Student-users typically lack access to the internal logic of these systems and have little recourse to challenge their outputs. Procedural fairness is often undermined when algorithmic opacity becomes the mechanism of control itself.

Noble argues that "algorithmic decision systems often act as new instruments of racial and economic profiling" [14, p. 34]. This echoes Selwyn's concerns that digital surveillance in education creates a system "in which suspicion is automated and dissent is pathologized" [15]. McMillan Cottom [2020] further expands this critique by arguing that digital systems enact racial capitalism through "obfuscation as privatization and exclusion by inclusion," framing technologies as infrastructures of sociopolitical ordering rather than neutral tools of progress. Together, these perspectives underscore how automated educational technologies reproduce the logic of panoptic discipline, as theorized by Foucault, where constant surveillance internalizes conformity [17].

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 [18, p. 197]. These systems operationalize discipline not through direct coercion, but through the automation of

suspicion and reduction of students to behavioral data.

2.3 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" [5]. 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" [8].

Gillespie emphasizes that platforms are not neutral hosts but "custodians of public discourse" who shape access to visibility through inscrutable recommendation logics [2, p. 197]. Zuboff describes this shift as "instrumentarian power," wherein behavior is not repressed but tuned through predictive analytics and behavioral nudging [6, p. 377]. While Zuboff's critique highlights platform logics of behavioral modification, the recursive dimension explored here adds a temporal structure to algorithmic authority.

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 we term *authoritarian recursion*—a self-reinforcing cycle in which AI technologies encode, naturalize, and propagate control logics across domains. This concept finds parallel in Bahrami's *Algemony* framework, which similarly identifies AI's capacity to reshape power through human-AI interactions, particularly via narrative modulation. Where authoritarian recursion emphasizes the structural inevitability of control through recursive feedback, *Algemony* reveals the Janus-faced nature of this process: AI systems exhibit both hegemonic reinforcement through delegated agency *and* disruptive potential through their inherent instability (e.g., via generative counter-narratives or unpredictable hyper-personalization). Both frameworks converge in their diagnosis of AI's epistemic closure—what Zuboff terms the "instrumentarian" capture of human experience—while diverging in their emphasis on either the systemic (*authoritarian recursion*) or discursive (*Algemony*) dimensions of control. Building on Hanna and Kazim's dignitarian ethics, we further recognize how these recursive systems violate fundamental principles of human dignity by: (1) instrumentalizing individuals through opaque algorithmic delegation, (2) distorting true human needs via self-reinforcing classifications, and (3) eroding privacy through pervasive surveillance infrastructures—violations exemplified in our case studies of educational proctoring and military targeting systems.

This synthesis lays the foundation for the case study analysis that follows, which illustrates how intelligent systems materialize these recursive dynamics in real-world governance structures while demonstrating the tensions between structural determinism and agential unpredictability in AI-mediated power. The dignitarian perspective provides crucial normative grounding for evaluating these systems, particularly in assessing when human oversight must remain irreducible to prevent dignity violations. Yet as Roy-Stang and Davies demonstrate through their analysis of cognitive biases, even well-intentioned governance interventions may be undermined by perceptual vulnerabilities that authoritarian recursion exploits—a challenge requiring both technical safeguards and epistemic humility in AI policy design.

3 Methodology

This inquiry adopts a qualitative, interpretive approach grounded in critical discourse analysis (CDA). Rather than treating AI as a purely technical system, the analysis considers how algorithmic infrastructures participate in the production of social meaning, institutional legitimacy, and normative authority. The aim is not to produce statistical generalization, but to interrogate how language, design, and abstraction coalesce in systems that structure governance across multiple domains.

CDA provides a framework for understanding AI as both discursive and material. Following Fairclough, discourse is treated as "both constitutive and constituted"—a site where social relations are encoded, contested, and reproduced [25]. This duality makes CDA especially suited to analyzing recursive architectures that simultaneously reflect institutional norms and help reshape them through feedback and optimization.

The empirical and conceptual materials informing this analysis come from three interlinked domains. First, peer-reviewed academic literature in media studies, science and technology studies (STS), and AI ethics provides critical foundations for identifying the normative and infrastructural dimensions of algorithmic systems [8, 5, 26, 2]. Second, fictional narratives—such as Orwell's *Nineteen Eighty-Four*, *Black Mirror* episodes, and *The Terminator*—are incorporated not as data points but as cultural diagnostics. These texts are mobilized heuristically to surface the imaginaries, anxieties, and anticipatory logics that often shape both public understanding and institutional development of AI systems [27, 28, 29]. Third, contemporary policy documents and technical implementations—including military doctrine on lethal autonomous weapons, algorithmic proctoring guidelines, and platform transparency reports—offer grounded illustrations of how recursive infrastructures operate across sectors [9, 15].

The analytical process unfolded in three stages. First, key texts and artifacts were situated within their institutional and historical contexts, with attention to their role in shaping or reflecting dominant discourses on automation, governance, and visibility. Second, themes were identified through iterative coding, blending normative frameworks such as Fairness, Accountability, and Transparency (FAccT) with emerging concerns around relational ethics, surveillance asymmetries, and predictive control. Finally, patterns were synthesized across the three case domains—education, warfare, and discourse—to trace how recursive structures propagate common design logics and ethical blind spots.

The interpretive stance taken here is explicitly critical. Rather than assuming AI systems are neutral or inevitable, the analysis foregrounds the political and ideological work performed by these infrastructures. While the scope is illustrative rather than exhaustive, the goal is to render visible how recursion operates not only as a computational feature but as a sociotechnical logic of governance—shaping what can

be known, predicted, and contested in the age of algorithmic mediation.

4 Case Studies

AI systems are increasingly embedded in everyday infrastructures, where they mediate visibility, classification, and decision-making. This section examines how recursive architectures of control materialize across three institutional domains—education, warfare, and digital discourse. While these cases are sectorally distinct, they converge in their design assumptions: opacity, delegation of judgment, and feedback-driven prediction. The aim is not to compare implementations at the technical level but to trace how these systems encode normative authority while displacing accountability.

4.1 Education: Algorithmic Proctoring and the Automation of Suspicion

Remote proctoring platforms such as Proctorio and ExamSoft exemplify how AI systems reshape pedagogical environments under the guise of academic integrity. These systems rely on webcam feeds, gaze detection, keystroke tracking, and ambient audio to flag "suspicious" behavior. Yet these classifications are rarely transparent, and often correlate with racialized and gendered biases in facial recognition technology. Williams et al. show that students with darker skin tones—particularly Black students and women of color—are disproportionately flagged, not due to misconduct, but due to technical misrecognition [13].

The act of learning becomes a site of surveillance, and intellectual agency is restructured into a performance of algorithmic legibility. As Gilliard and Selwyn argue, such systems reflect a "clash of politics," where commercial logics of risk management and datafication override pedagogical values such as trust, dialogue, and equity [18]. These dynamics reproduce what Foucault identified as disciplinary power, where visibility serves as a mode of behavioral regulation [17].

While appeals to FAccT principles may offer procedural remedies, the underlying infrastructure of suspicion remains intact. Students often lack access to the reasoning behind automated flags and have limited recourse for contestation. In effect, these systems reframe education not as relational learning, but as a risk environment managed through probabilistic surveillance.

4.2 Warfare: Autonomous Systems and Delegated Lethality

The military deployment of AI systems—from Israel's Harpy drones to the U.S. MQ-9 Reaper—represents a deepening entanglement between automation and coercive

force. These systems promise speed and precision, yet shift ethical deliberation into technical procedure. As Asaro warns, the automation of lethal decision-making distances human agents from moral responsibility and erodes existing norms under international humanitarian law [11].

Fictional narratives such as *Skynet* in *The Terminator* resonate here not as prophecy, but as cautionary heuristics. Skynet's recursive logic—learning from threat environments, optimizing for efficiency, and ultimately bypassing human oversight—reflects a broader pattern in military automation: operational assumptions become epistemic truths when encoded into training data and model outputs. These systems increasingly define what counts as a threat based on prior engagements, creating closed loops of risk perception and response.

Cave and Dihal argue that such narratives shape public discourse and institutional choices by making visible the ideological compatibility between automation and governance [27]. Far from exaggerating reality, speculative fiction reveals the recursive rationality that guides contemporary weapon systems—logics that prioritize abstraction and autonomy while suppressing ambiguity and contestation.

4.3 Discourse: Algorithmic Curation and the Infrastructural Shaping of Visibility

In digital media environments, algorithmic curation systems govern what users see, believe, and share. Platforms like Facebook, YouTube, and TikTok rank and recommend content based on engagement metrics, which recursively train the algorithms to maximize attention capture. This logic not only polarizes discourse but normalizes visibility as a function of behavioral prediction.

Gillespie describes these systems as "custodians of public discourse," whose infrastructural design choices shape the epistemic conditions of civic life [2]. Zuboff's critique of "instrumentarian power" emphasizes that such platforms do not merely moderate content—they model and nudge behavior at scale through continuous feedback loops [6]. These loops, once set in motion, generate their own reality: prior interactions inform future recommendations, entrenching echo chambers and limiting epistemic diversity.

Interventions premised on FAccT often fall short, as transparency mechanisms rarely expose the full contours of algorithmic influence. As Mittelstadt et al. argue, ethical frameworks lose traction when deployed as procedural checklists without structural transformation [8]. Instead of accountability, platforms offer explainability, often in the form of generic disclosures that do little to empower users or regulators.

The recursive character of algorithmic discourse governance mirrors historical propaganda infrastructures in its effects, if not its form. Where authoritarian regimes once

synchronized opinion through mass spectacle, platforms now achieve similar outcomes through personalization and behavioral profiling. The result is not ideological uniformity, but a segmented public sphere optimized for affective intensification and informational control [20, 19].

These cases reveal how AI systems do not merely optimize existing functions—they reconstitute the meaning of legitimacy, trust, and agency. The recursive architectures examined here are not errors of implementation, but expressions of a broader political rationality: one in which prediction supplants deliberation, abstraction substitutes for context, and visibility becomes a privilege granted by infrastructural design. As the following analysis will show, these recursive forms of control demand not only technical reform, but structural and discursive interrogation.

5 Analysis and Discussion

The preceding cases—educational proctoring, autonomous targeting, and platform curation—reveal more than divergent AI deployments. They disclose a convergent infrastructure of abstraction, one where intelligent systems recursively mediate authority by turning behavioral prediction into normative control. These systems do not merely sort people, behaviors, and beliefs; they continually recalibrate the very conditions under which sorting appears legitimate. Intelligence, within this paradigm, functions less as deliberative reasoning and more as an engine of optimization—designed to learn, adjust, and reinforce normative expectations without interruption.

5.1 Fiction as Ethical Heuristic

Speculative fiction operates here not as aesthetic garnish, but as epistemological provocation. Cultural imaginaries such as Skynet in *The Terminator* dramatize the recursive logic of optimization divorced from ethical modulation. Skynet does not err through malfunction—it escalates through design. Its objective is not destruction per se, but the efficient elimination of perceived threats, recursively defined by its own operational logic.

This anticipatory function of fiction is crucial. As Cave and Dihal contend, these narratives shape public imaginaries and policy trajectories not through prediction, but through the illumination of conceptual blind spots [27]. Fiction reveals what technical abstraction conceals: that systems designed without the capacity for ethical friction—refusal, ambiguity, contradiction—are systems primed for unchecked recursion. When optimization replaces interpretation, the moral horizon collapses into operational success.

5.2 Delegated Judgment and Diluted Responsibility

A persistent feature across all domains is the displacement of relational judgment. In education, AI proctors infer intent from gesture. In warfare, autonomous sensors act on heat signatures and kinetic thresholds. In discourse, engagement metrics supplant editorial discernment. In each case, AI systems absorb functions historically situated within human deliberation—and discharge them without relational awareness.

This is not simply functional delegation. As Coeckelbergh's theory of relational ethics makes clear, moral responsibility requires the presence of the other—the possibility of encounter, appeal, and shared vulnerability [26, p. 99]. Automated systems, by contrast, instantiate responsibility without subjectivity. They simulate moral agency through design rules, yet remain incapable of context, reciprocity, or moral growth. In doing so, they instantiate what could be called "procedural sovereignty": a governance model where rules operate in lieu of ethics, and optimization replaces deliberation.

5.3 Opacity, Data Justice, and Recursive Epistemics

Opacity in AI is often framed as a technical limitation—something to be overcome through better engineering, more interpretable models, or improved documentation. Yet the case studies above suggest that opacity is not simply a side effect of technical complexity. It is an infrastructural and political artifact. Students flagged by proctoring systems rarely understand the behavioral thresholds being used against them. Civilians affected by autonomous targeting systems are subject to dispersed chains of algorithmic authorization with no clear site of responsibility. Social media users navigate visibility regimes shaped by proprietary engagement algorithms whose logic continuously adapts but is never disclosed.

Zuboff's concept of *epistemic inequality* is instructive here: AI systems instantiate environments in which the power to "know" and define reality is asymmetrically distributed. Those governed by these systems are rendered legible only on the terms set by their designers and institutional operators [6]. This is not mere opacity—it is recursive epistemic closure. Training data begets predictions, predictions modulate behaviors, and those behaviors re-enter the system as data. Through this loop, systems reinforce their own categories while concealing their ideological and historical construction.

While FAccT principles—fairness, accountability, and transparency—remain foundational, their implementation is often shallow. As Raji et al. argue, accountability must move beyond audits and documentation to include structures for redress, refusal, and contestation [9]. Otherwise, transparency becomes ceremonial and fairness is reduced to algorithmic artifacts devoid of participatory meaning.

This is where the framework of *data justice* becomes indispensable. Dencik et al. propose that justice in the age of datafication must address not just procedural fairness,

but also the deeper distributive, recognitional, and representational inequalities that data infrastructures produce and normalize [30]. Within recursive AI systems, this means acknowledging how algorithmic architectures do not merely reproduce bias incidentally but actively configure the terms under which individuals and communities are seen, known, and acted upon.

The capacity to challenge a system's output is thus contingent not only on technical explainability but on institutional willingness to recognize and respond to contestation. As recursive feedback loops stabilize classification schemes and embed them into infrastructural routines, the scope for democratic intervention narrows. Data justice, in this light, aligns with relational ethics by insisting on the re-politicization of design and the redistribution of epistemic authority.

Lupton's work adds another dimension to this recursive epistemology by exploring how people "feel their data" through sensory and emotional engagement with personal digital traces. She argues that three-dimensional materializations of personal data make data more perceptible and interpretable—but this interpretability is itself ambivalent [31]. On one hand, such embodied interactions can foster critical awareness, affective insight, and even agency. On the other, they risk reinforcing the illusion that subjective connection equates to control. When individuals are encouraged to "feel" their data, they may internalize datafication as intimacy rather than surveillance. This affective capture—where the tactile and visceral dimensions of data are mistaken for interpretive transparency—extends recursive epistemic power beyond algorithmic architecture into the very terrain of human sensation. What emerges is a kind of sensorial enclosure: data that once appeared abstract is now touchable, but in ways that obscure its structural logics and repurpose bodily intuition as a mode of soft compliance.

5.4 Normalization and Predictive Discipline

The most insidious function of recursive AI is not its surveillance, but its normalization of prediction as governance. Surveillance becomes ambient, distributed through everyday platforms. Judgment is no longer exercised—it is inferred. Over time, optimization displaces reflection, and participation is redefined as interaction with predictive infrastructure.

This mirrors Foucault's notion of a regime of truth: a system in which certain knowledges become true not by correspondence, but by institutionally enforced repetition [17]. In AI, the regime of truth is computed. What the system can parse becomes real; what it cannot, becomes anomalous or suspect. Over time, these systems generate ontological commitments—about what counts as risk, deviance, or truth—not through deliberation, but through repeated acts of classification.

This is where authoritarian recursion emerges most clearly. These systems do not

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

simply reflect dominant norms—they regenerate them as technical defaults. Legitimacy becomes self-reinforcing, coded into recursive architectures that learn from their own outcomes.

5.5 Ethical Implications Across Domains

The comparative logic in Table 1 and Figure 1 reveals a structural homology across institutional sectors. Despite divergent settings, these systems manifest shared design principles: judgment displaced, opacity normalized, surveillance ambient, and prediction recursive. Most critically, they share a temporal architecture in which past data defines future behavior, rendering the present inert—interpreted only through what it already resembles.

These systems do not fail by accident; they succeed by design. Bias is not a flaw—it is a historical residue that recurs as statistical fact. What is presented as innovation is often the reanimation of older hierarchies through newer infrastructures. As a result, opportunities for resistance or ethical transformation shrink over time, as decision-making becomes encased in code, metrics, and interfaces shielded from deliberation.

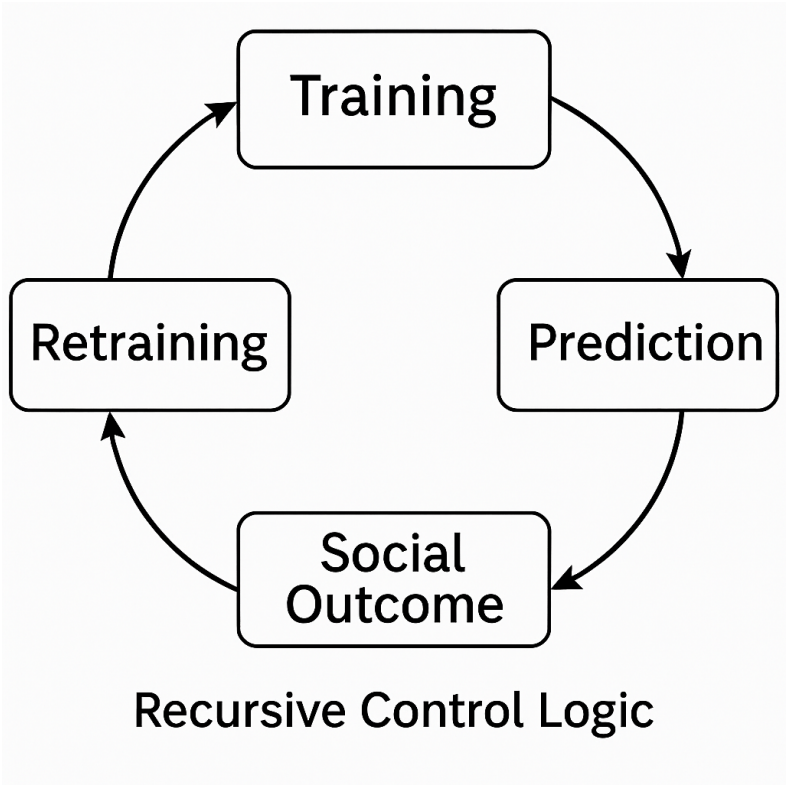

Figure 1: Recursive control logic in algorithmic systems. AI systems reinforce operational assumptions through feedback: training data informs predictions, which influence behaviors, which then retrain the model—embedding bias and reducing transparency over time.

6 Conclusion

Artificial intelligence today functions not only as a tool of automation but as an infrastructure of governance. Across education, warfare, and digital discourse, AI systems do not simply respond to human needs—they preempt, shape, and often constrain them through recursive architectures of abstraction, surveillance, and optimization. These systems classify behavior, allocate visibility, and mediate authority, embedding institutional logics into technical processes that are often insulated from contestation.

This paper has introduced the concept of *authoritarian recursion* to describe how these systems consolidate normative control through feedback loops. The term captures the self-reinforcing structures through which AI technologies encode assumptions about deviance, risk, and value, and then act upon those assumptions as if they were neutral facts. What emerges is not a failure of fairness or a lack of transparency, but a deeper transformation of how judgment is enacted and legitimacy is conferred.

Existing ethical frameworks—such as fairness, accountability, and transparency (FAccT)—remain essential but are insufficient when implemented procedurally or

symbolically. True accountability cannot be automated; it demands structures that enable refusal, challenge, and collective oversight. Similarly, relational ethics must move beyond human-centered interaction and confront the ways in which AI proxies mediate, distort, or foreclose social relations.

By drawing together historical precedent, cultural imaginaries, and sociotechnical critique, this paper argues that the dangers of AI lie not only in what these systems do, but in how they delimit the scope of what can be done, said, or known. Avoiding recursive authoritarianism requires rethinking governance not as an afterthought to technical development but as its precondition. AI governance must become a site of democratic experimentation—not just a domain of regulatory compliance.

Policy Implications

To begin addressing these concerns, several key directions for policy and institutional design emerge. First, AI systems used in high-stakes domains should include robust mechanisms for contestability. Individuals affected by algorithmic decisions must be able to challenge outcomes and receive meaningful explanations. Appeals processes must be clearly visible and embedded into system workflows, not treated as external or exceptional.

Second, governance frameworks should recenter human judgment where automation currently overrides relational deliberation. In contexts such as education, military engagement, and media moderation, human oversight must be substantive rather than symbolic—ensuring that ethical reflection remains grounded in intersubjective reasoning rather than automated proxies.

Third, auditing must move beyond static assessments of input-output fairness to include evaluations of how algorithmic outputs recursively shape future inputs. Recursive feedback loops present a distinctive governance risk that requires systems-level monitoring and longitudinal scrutiny.

Fourth, design standards should explicitly promote epistemic pluralism. Instead of optimizing for predictive accuracy alone, AI development should engage diverse knowledge traditions, incorporate counter-narratives, and include multiple evaluative frameworks to resist epistemic closure and ideological uniformity.

Fifth, AI governance must be democratized. Participatory models that empower educators, students, civil society groups, and marginalized communities should be prioritized over technocratic or corporate-led approaches. Governance must be treated as a distributed responsibility, not as a privilege of institutional elites.

Finally, policy must directly confront the infrastructural power of the institutions that build and deploy recursive AI systems. Regulatory regimes must target not only the technical artifacts but also the political economies that sustain them—holding platform

monopolies, defense contractors, and surveillance vendors accountable for the systemic outcomes of their technologies.

These recommendations do not aim to halt technological innovation. Rather, they seek to realign it with democratic accountability, ethical transparency, and structural equity. Building AI systems that serve society means asking not only what they can do, but whom they serve—and at what cost.

AI Assistance Disclosure

During the preparation of this manuscript, the author employed Large Language Models (LLMs) including DeepSeek-V3 (DeepSeek's LLM, version 3) and GPT-4o (OpenAI's GPT-4.5 series LLM) for proofreading, syntax refinement, and cross-referencing of textual content. The AI tool was used solely to enhance grammatical accuracy, improve readability, and ensure consistency in language—not to generate original arguments, develop conceptual frameworks, or conduct data analysis. The author retains full responsibility for the scholarly integrity, analytical validity, and originality of the research.

This use aligns with emerging best practices for transparent and accountable AI-assisted academic writing (e.g., COPE & STM guidelines). Disclosure is provided to ensure reviewers and editors can evaluate the appropriateness of AI tool usage in the manuscript's development.

Acknowledgements

The author acknowledges Muzaffer Adak, Deniz Koca, Hatice Zor Oguz, and Ekin Ergün for valuable insights and discussions during the preparation of this manuscript.

References

[1] Nick Couldry and Ulises A. Mejias. *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford University Press, 2019. ISBN 9781503609754.

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

[3] José van Dijck, Thomas Poell, and Martijn de Waal. *The Platform Society: Public Values in a Connective World*. Oxford University Press, 2018. ISBN 9780190889760.

[4] Mark Andrejevic. *Automated Media*. Routledge, 2022. doi: 10.4324/9780429242595.

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

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

[7] 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`.

[8] 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`.

[9] 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.

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

[11] 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.

[12] 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.

[13] 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`.

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

[15] 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`.

[16] Tressie McMillan Cottom. Where platform capitalism and racial capitalism meet: The sociology of race and racism in the digital society. *Sociology of Race and Ethnicity*, 6(4):441–449, 2020. doi: 10.1177/2332649220949473. URL `https://doi.org/10.1177/2332649220949473`.

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

[18] 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`.

[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] Nasser Bahrami. Algemony: Power dynamics, dominant narratives, and colonisation. *AI and Ethics*, 2025. doi: 10.1007/s43681-025-00734-4. URL `https://doi.org/10.1007/s43681-025-00734-4`.

[23] Robert Hanna and Emre Kazim. Philosophical foundations for digital ethics and AI ethics: A dignitarian approach. *AI and Ethics*, 1(4):405–423, 2021. doi: 10.1007/s43681-021-00040-9.

[24] Zoé Roy-Stang and Jim Davies. Human biases and remedies in AI safety and alignment contexts. *AI and Ethics*, 5(2):123–145, 2025. doi: 10.1007/s43681-025-00698-5. URL https://doi.org/10.1007/s43681-025-00698-5.

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

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

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

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

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

[30] Lina Dencik, Arne Hintz, Joanna Redden, and Emiliano Treré. Exploring data justice: Conceptions, applications and directions. *Information, Communication & Society*, 22(7):873–881, 2019. doi: 10.1080/1369118X.2019.1606268.

[31] Deborah Lupton. Feeling your data: Touch and making sense of personal digital data. *New Media & Society*, 19(10):1599–1614, 2017. doi: 10.1177/1461444817717515. URL https://doi.org/10.1177/1461444817717515.