AI and Social Aspects

Weekly Analysis — https://ainews.social

The Architecture of Power in AI Discourse: Who Speaks, Who Suffers, and Who Profits

When a student in Lawrence, Kansas is flagged by AI surveillance software for making a joke about bringing a "bomb" Pop-Tart to school, we witness more than a technical glitch. We see the crystallization of power relations that define contemporary AI discourse: those who design speak, those who deploy decide, and those who suffer remain silent. The discourse around artificial intelligence in education and society reveals a profound asymmetry—between who shapes the conversation and who bears its consequences, between who profits from AI systems and who labors invisibly to build them, between whose expertise counts and whose experience is dismissed. As documented in [10], students subjected to algorithmic surveillance have limited recourse beyond lawsuits, while the companies building these systems shape policy discussions about "safety" and "protection."

[10] Lawrence school district sued in federal court for use of AI-powered ...

This essay interrogates the power structures embedded in AI discourse, revealing how the conversation itself—who gets to speak, what counts as expertise, which problems matter—reproduces and amplifies existing inequalities. The evidence reveals a discourse dominated by corporate perspectives focused on productivity and optimization, while the voices of those most harmed by algorithmic systems—data workers in the Philippines, students under surveillance, job seekers rejected by biased hiring algorithms—remain systematically excluded. When 39.5% of analyzed articles identify ethical failures in AI systems, we must ask: who defines these failures, who documents them, and crucially, who escapes accountability for creating them?

The patterns are stark. Western institutions and Global North perspectives overwhelmingly shape AI discourse, as revealed by analyses showing U.S.-centric and Eurocentric viewpoints dominating while Global South experiences remain marginalized. Technical solutionism prevails, with researchers proposing algorithmic fixes to problems that are fundamentally about power, resources, and justice. Most tellingly, those with the least power in the AI ecosystem—the workers who label data, the students under surveillance, the communities experiencing algorithmic bias—are precisely those whose perspectives appear least

in the discourse that ostensibly addresses their concerns.

The Architecture of Silence: Whose Voices Shape AI's Future

The structure of AI discourse reveals itself first through its silences. While researchers, policymakers, and technologists debate ethics and governance, entire categories of stakeholders remain systematically excluded from these conversations. The most glaring absence involves those who perform the hidden labor that makes AI possible. As documented in [5], workers in the Philippines endure grueling conditions labeling data for AI systems, yet their perspectives on AI development remain entirely absent from policy discussions about "responsible AI." These workers, earning minimal wages while viewing traumatic content to train AI models, represent the human infrastructure of artificial intelligence, yet they have no seat at tables where AI governance is discussed.

[5] Aux Philippines, le calvaire des petites mains de l'intelligence ...

This exclusion is not accidental but structural. The discourse privileges certain forms of expertise—technical, managerial, regulatory—while systematically devaluing experiential knowledge. When [7] reveals increasing skepticism among developers about AI reliability, their doubts carry weight in shaping industry practices. But when students report false alarms from AI surveillance systems, as detailed in [14], their experiences are treated as anecdotal rather than systemic evidence of harm.

[7] Developers Lean on AI More, But Report Growing Doubts ...

[14] Students have been called to the office for AI surveillance false alarms

The architecture of silence extends to entire regions. Research consistently shows Western institutions dominating AI discourse, with consequences far beyond academic publishing. As [17] demonstrates, AI educational tools designed in the Global North embed cultural biases that erase African languages, histories, and ways of knowing. This isn't merely oversight—it's a form of data colonialism where AI systems trained on Western data reproduce and amplify existing global inequalities. African educators implementing these tools face an impossible choice: adopt systems that undermine local cultures or be excluded from the "AI revolution" entirely.

[17] The cultural cost of AI in Africa's education systems - UNESCO

Even when Global South perspectives appear in discourse, they often do so through Northern intermediaries. UNESCO reports and World Bank studies speak about African education rather than amplifying African educators' voices. This ventriloquism maintains existing power structures while creating an illusion of inclusion. The result is a discourse that discusses the "global" impact of AI while systematically excluding most of the globe from meaningful participation in shaping AI's development.

When Harm Becomes Data: The Transformation of Suffering into Research

The AI discourse exhibits a troubling pattern: extensive documentation of harm coupled with minimal accountability for those causing it. When [2] catalogs discriminatory hiring algorithms, or when [3] systematically reviews educational AI's discriminatory impacts, these studies transform lived experiences of exclusion into data points for academic analysis. The violence of algorithmic bias becomes sanitized through research methodologies, converted into conference presentations and policy recommendations while those experiencing discrimination continue to be rejected by hiring systems or miseducated by biased tools.

[2] AI Hiring Bias: Real Cases, Legal Consequences, and Prevention

[3] Algorithmic Bias in Education

This transformation serves a specific function in maintaining power relations. By converting harm into an object of study, the discourse shifts focus from accountability to analysis. Companies deploying biased systems evade responsibility while researchers build careers documenting bias. The extensive evidence of AI surveillance's negative impacts on students, detailed in [13], generates academic papers and policy briefs but rarely leads to these systems being removed from schools. Instead, the harm becomes productive—generating research funding, conference panels, and consulting opportunities for those positioned to study it.

[13] Schools are using AI to spy on students and some are getting arrested ...

The 40.8% of articles identifying ethical failures in AI systems represents more than a statistic—it reveals how the discourse metabolizes harm. Each documented failure becomes evidence for the next grant proposal, the next research project, the next framework for "responsible AI." Yet this accumulation of evidence rarely translates into material changes for those harmed. Students continue to be surveilled, job seekers continue to face biased algorithms, and Global South communities continue to experience data extraction. The discourse's productivity depends on the continuation of harm, creating perverse incentives where documenting problems becomes more valuable than solving them.

Consider how [6] provides a "harmonized framework" for understanding bias in educational AI. While such frameworks offer analytical clarity, they also abstract away the concrete experiences of students whose educational opportunities are curtailed by biased systems. The framework becomes a technical artifact that researchers can cite and build upon, while the students it ostensibly serves remain objects of study rather than agents of change. This dynamic reveals how power operates through discourse—by transforming political problems into technical challenges requiring expert management rather than

[6] FairAIED: Navigating Fairness, Bias, and Ethics in Educational AI ...

fundamental restructuring.

The Governance Mirage: How Regulation Talk Obscures Power

The proliferation of governance frameworks and regulatory discussions masks a deeper truth about power in AI discourse. While 39.2% of articles frame AI as a "governance challenge," this framing itself serves specific interests. As [18] demonstrates through its comprehensive review, the focus on governance often assumes AI adoption is inevitable and beneficial, requiring only proper management. This assumption removes fundamental questions about whether AI systems should be deployed in educational settings from democratic deliberation.

[18] Towards responsible artificial intelligence in education: a systematic ...

The governance discourse operates through a peculiar logic. It acknowledges risks—bias, surveillance, privacy violations—while simultaneously legitimizing the systems that create these risks. When [1] outlines regulatory trends, it frames governance as a technical challenge for experts rather than a political question for communities. This expertification of governance excludes those most affected by AI systems from decisions about their deployment. Students subjected to surveillance, workers replaced by algorithms, and communities experiencing algorithmic discrimination are positioned as beneficiaries of better governance rather than participants in defining what governance means.

[1] 11 things AI experts are watching for in 2026

International governance discussions reveal these power dynamics starkly. The European Union's AI Act, discussed in [7], represents years of negotiation among policymakers, industry lobbyists, and technical experts. Conspicuously absent from these negotiations were representatives of communities experiencing algorithmic discrimination. The resulting framework, while addressing some forms of bias, embeds assumptions about AI's inevitability and legitimacy that preclude more fundamental challenges to algorithmic power.

[7] Gaps and policies in AI- and algorithm-driven discrimination in Europe

Even critical governance discussions often reinforce existing power structures. When [16] argues for moving beyond technical solutions, it proposes "dialogic reflection" among educators and administrators. While valuable, this approach still locates power within institutional structures rather than challenging the fundamental asymmetry between those who control AI systems and those subjected to them. The governance discourse, even in its critical modes, tends to manage rather than challenge algorithmic power.

[16] Technical fixes alone cannot solve AI bias in education

Laboring in Digital Shadows: The Hidden Workers of

AI

The most profound power asymmetry in AI discourse involves those whose labor makes these systems possible yet who remain invisible in discussions about AI's future. Data annotation workers, content moderators, and the vast workforce of the Global South who train AI models exist in what might be called digital shadows—essential yet unseen, productive yet powerless. [12] reveals how workers in the Philippines face exploitative conditions while performing the psychological labor of making AI systems function, viewing disturbing content for minimal wages while being excluded from any meaningful participation in shaping the industry that depends on their work.

[12] Philippines - Intelligence artificielle (IA): Les forçats de l'IA ...

This invisibility is actively produced through discourse. When [4] analyzes AI's economic impact, it focuses on productivity gains and job displacement in formal sectors while ignoring the informal economy of data work that underpins AI development. The report's "economic primitives" framework captures certain forms of value while rendering others invisible. This selective visibility ensures that discussions about AI's economic impact center on relatively privileged workers in the Global North while ignoring the exploitative conditions faced by data workers in the Global South.

[4] Anthropic Economic Index report: Economic primitives

The power dynamics become even starker when examining how AI companies structure these labor relations. Workers are typically classified as independent contractors, denied basic labor protections, and prevented from organizing collectively. As documented in [15], these "taskers" face algorithmic management that surveils their productivity while denying them any agency in defining their working conditions. The same companies promoting "human-centered AI" and "AI for good" depend on labor practices that deny workers their humanity.

[15] Taskers: los trabajadores precarios digitales que dan vida a la IA

The discourse's treatment of AI labor reveals its ideological function. By focusing on automation's impact on traditional employment, as in [8], the conversation obscures the new forms of labor AI creates—forms characterized by invisibility, precarity, and exploitation. This isn't merely an oversight but a structural feature of how AI discourse operates, maintaining the fiction that AI systems emerge from pure computation rather than human labor. The workers who make AI possible are excluded from conversations about AI governance, ethics, and futures precisely because acknowledging their existence would reveal uncomfortable truths about the industry's dependence on digital sweatshops.

[8] L'IA remplacera-t-elle des emplois ? Un rapport d'Anthropic ...

The Geography of AI Power: Center, Periphery, and Algorithmic Colonialism

The spatial distribution of power in AI discourse maps directly onto existing global inequalities, creating what scholars increasingly recognize as algorithmic colonialism. This isn't merely about where AI companies are headquartered or where research is published—it's about whose data is extracted, whose cultures are erased, and whose futures are foreclosed by AI systems designed elsewhere. [11] provides empirical evidence that large language models systematically favor Western perspectives, languages, and cultural assumptions, literally encoding global power relations into their outputs.

[11] New study finds that ChatGPT amplifies global inequalities

The African context illuminates these dynamics with particular clarity. When [17] documents how AI educational tools erase African languages and knowledge systems, it reveals algorithmic colonialism in action. African students learn from AI systems trained on Western data, embedding Western assumptions about knowledge, learning, and culture. This isn't a bug but a feature of how AI development is currently structured—with data extraction from the Global South feeding models designed in the Global North that then shape how Southern populations access education and information.

[17] The cultural cost of AI in Africa's education systems - UNESCO

The discourse itself reproduces these colonial relations. Research about AI in Africa, as seen in comprehensive reports, typically emerges from Northern institutions studying Southern contexts. Even well-intentioned efforts to address bias often reinforce center-periphery dynamics. When [20] reveals how ChatGPT's rankings favor wealthy Western nations, the study emerges from Oxford rather than from institutions in the regions being misrepresented. This pattern—Northern institutions documenting Southern disadvantage—maintains existing hierarchies even while critiquing them.

[20] « The Silicon Gaze » : Les classements de ChatGPT ...

Latin American responses to AI colonialism offer glimpses of alternative possibilities. [19] documents how journalists in the region develop community-centered approaches to AI that prioritize local knowledge and needs. Yet these initiatives struggle for visibility in a discourse dominated by English-language publications and Western frameworks. The geography of AI power ensures that alternatives to dominant models remain marginalized, discussed in regional forums rather than shaping global AI governance.

[19] When journalism reclaims agency: Latin America's AI answer

Conclusion: Redistributing Power in the Age of Algorithms

The analysis reveals a discourse structured by and for existing power relations. Those who profit from AI systems—technology companies, Global North institutions, research establishments—dominate conversations about AI's development and deployment. Those who

suffer AI's harms—surveilled students, exploited workers, discriminated job seekers, marginalized communities—appear in discourse primarily as objects of study rather than subjects with agency. This asymmetry isn't incidental but fundamental to how AI discourse operates, transforming political questions about power into technical questions about optimization.

The evidence points toward necessary interventions. First, discourse itself must be democratized, creating spaces where those experiencing AI's impacts can speak from their expertise rather than being spoken about. Second, the hidden labor of AI must be made visible and valued, recognizing data workers as essential stakeholders in AI governance. Third, alternatives to dominant AI models emerging from the Global South need amplification and resources, challenging the assumption that AI development must follow patterns established in Silicon Valley.

As [9] suggests, we stand at a crossroads between technological rupture and pedagogical continuity. But this framing itself obscures the more fundamental choice: between reproducing existing power relations through algorithmic systems and reimagining how power might be distributed in an age of artificial intelligence. The discourse's current structure—dominated by ethical failure narratives, governance frameworks, and technical solutionism—serves those already powerful while offering little to those seeking liberation from algorithmic control.

[9] L'intelligence artificielle en éducation : entre rupture ...

The path forward requires more than better ethics frameworks or more inclusive governance. It demands fundamental questioning of who benefits from AI systems, who controls their development, and who bears their risks. Until the discourse includes those it currently silences—the data workers, the surveilled, the marginalized—it will continue to reproduce the very power asymmetries it occasionally documents but rarely challenges. The question is not how to govern AI better but how to redistribute the power that AI currently concentrates, ensuring that those who suffer its harms have equal say in shaping its futures.

References

1. 11 things AI experts are watching for in 2026
2. AI Hiring Bias: Real Cases, Legal Consequences, and Prevention
3. Algorithmic Bias in Education
4. Anthropic Economic Index report: Economic primitives

5. Aux Philippines, le calvaire des petites mains de l'intelligence ...
6. FairAIED: Navigating Fairness, Bias, and Ethics in Educational AI ...
7. Gaps and policies in AI- and algorithm-driven discrimination in Europe
8. L'IA remplacera-t-elle des emplois ? Un rapport d'Anthropic ...
9. L'intelligence artificielle en éducation : entre rupture ...
10. Lawrence school district sued in federal court for use of AI-powered ...
11. New study finds that ChatGPT amplifies global inequalities
12. Philippines - Intelligence artificielle (IA): Les forçats de l'IA ...
13. Schools are using AI to spy on students and some are getting arrested ...
14. Students have been called to the office for AI surveillance false alarms
15. Taskers: los trabajadores precarios digitales que dan vida a la IA
16. Technical fixes alone cannot solve AI bias in education
17. The cultural cost of AI in Africa's education systems - UNESCO
18. Towards responsible artificial intelligence in education: a systematic ...
19. When journalism reclaims agency: Latin America's AI answer
20. « The Silicon Gaze » : Les classements de ChatGPT ...