

AI and Social Aspects

Weekly Analysis — <https://ainews.social>

When artificial intelligence enters a classroom, a workplace, or a welfare office, it doesn't arrive alone. It brings with it a constellation of assumptions, biases, and power relations that are rarely questioned—because the people with the most at stake rarely get to shape the conversation. The discourse around AI and its social impacts reveals a troubling truth: those who control the narrative about AI's risks and benefits are rarely those who bear its consequences.

Consider the stark asymmetry revealed in recent evidence: while 41% of AI implementations result in ethical failures, the solutions proposed overwhelmingly come from the same institutions and regions that created these problems. As [15] documents, Western AI tools dominate African educational contexts, eroding local knowledge systems while the discourse about "responsible AI" remains firmly centered in Silicon Valley boardrooms and European regulatory chambers.

This essay interrogates the power relations embedded in AI discourse, revealing how certain voices dominate while others are systematically silenced. Through critical analysis of who speaks, who is spoken about, and who remains unheard, we can begin to understand why AI's social impacts remain so persistently harmful despite endless governance initiatives and ethical frameworks. The question isn't just what AI does to society—it's who gets to decide what counts as harm, what counts as a solution, and whose interests shape the future of automated decision-making.

The Architecture of Inevitability

The most insidious form of power in AI discourse is the assumption of its inevitability. Across institutional pronouncements, policy documents, and educational initiatives, AI is presented not as a choice to be debated but as a force of nature to be managed. This framing strips agency from communities, educators, and citizens who might otherwise question whether AI should be deployed at all.

The business school perspective exemplifies this manufactured inevitability. [19] doesn't ask whether future leaders should work with AI, but assumes they must, focusing entirely on how to prepare

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

[19] Comment l'Insead veut former les futurs dirigeants à travailler avec l'IA

them for this predetermined future. Similarly, legal frameworks adapt existing structures rather than questioning fundamental assumptions, as seen in [19], which treats AI's presence in legal proceedings as a given rather than a choice.

This discourse of inevitability serves powerful interests. By framing AI as unstoppable, tech companies avoid debates about whether their products should exist. Institutions avoid hard questions about alternatives. Most critically, it shifts the burden of adaptation onto individuals and communities least equipped to resist. When educators must learn to "work with" surveillance systems that generate false positives, as documented in [14], the question of whether schools should surveil students at all disappears from view.

The power to define AI as inevitable is the power to shape society without democratic input. It transforms political questions about values, priorities, and human dignity into technical questions about implementation and mitigation. This sleight of hand reveals the first major asymmetry in AI discourse: those who benefit from AI's expansion control the narrative about its necessity, while those who bear its costs are told only to adapt.

The Governance Fixation

The overwhelming focus on governance and regulation in AI discourse—with governance mentioned in 911 articles and regulation in 589—reveals another form of power: the capture of critical discourse by institutional actors. While communities experience AI's impacts directly, the conversation about these impacts happens in regulatory chambers, academic conferences, and corporate compliance departments.

[13] illustrates this dynamic perfectly. The tension between innovation promotion and citizen protection plays out in legislative halls, with tech companies and regulators as the primary actors. The citizens whose data will be harvested, whose opportunities will be algorithmically determined, appear only as abstract concerns to be "balanced" against economic interests.

This governance fixation serves multiple functions for powerful actors. First, it positions them as responsible stewards addressing AI's challenges, even as they profit from its expansion. Second, it channels resistance into approved pathways—comment periods, stakeholder consultations, impact assessments—that rarely challenge fundamental premises. Third, it creates an illusion of action that obscures the lack of real change. Despite hundreds of governance frameworks, algorithmic

[19] As AI-generated fake content mars legal cases, states want guardrails

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

[13] South Korea's 'world-first' AI laws face pushback amid bid to become leading tech power

mic bias persists, as [8] demonstrates with its finding that widely-used educational algorithms systematically disadvantage Black and Hispanic students.

The French Défenseur des droits report [9] exposes how algorithmic governance in public services creates new forms of discrimination while claiming to enhance fairness. The report's detailed analysis of welfare algorithms shows how governance discourse focuses on transparency and accountability while the systems themselves entrench existing inequalities.

Most tellingly, the governance fixation reveals whose concerns matter in AI discourse. Corporate liability, regulatory compliance, and innovation ecosystems dominate discussions, while the lived experiences of those subjected to algorithmic decisions remain peripheral. This isn't accidental—it reflects who has the power to convene meetings, fund research, and shape policy agendas.

Algorithmic Harms and Geographic Power

Perhaps nowhere is the power asymmetry in AI discourse more evident than in the geographic distribution of who documents harms versus who experiences them. The Global North dominates the production of knowledge about AI bias, while the Global South disproportionately bears its consequences. This isn't merely an academic concern—it shapes which harms get recognized, which solutions get funded, and whose interests drive the global AI agenda.

[2] reveals how "data colonialism" extends beyond resource extraction to the very framing of AI problems and solutions. When Western institutions define bias primarily in terms legible to their own regulatory frameworks, they miss forms of harm that don't fit these categories. The report documents how Global South workers perform the hidden labor of training AI systems—labeling data, moderating content, correcting errors—while remaining invisible in discussions about AI's future.

The harm documentation itself reflects power relations. [10] provides detailed evidence of how these systems perpetuate racial injustice, yet the communities most affected by over-policing rarely lead the research or policy conversations about solutions. Similarly, [19] exposes how medical AI replicates gender disparities, but the discourse remains dominated by technical fixes rather than addressing the systemic exclusion of women from medical research.

This geographic and demographic power imbalance has material

[8] PDF Are Algorithms Biased in Education? Exploring Racial Bias in Predicting ...

[9] PDF Rapport algorithmes, systèmes d IA et services publics : quels droits ...

[2] Assessing the AI Moment from a Global South Perspective

[10] Predictive policing algorithms are racist. They need to be dismantled.

[19] Gender bias in AI health advice: when innovation learns from incomplete data

consequences. [16] argues that bias emerges not just from algorithms but from the interaction patterns between humans and AI systems. Yet the humans studied, the interaction patterns deemed worthy of research, and the solutions proposed all reflect the priorities of well-resourced institutions in wealthy nations.

The exploitation goes deeper than bias in outputs—it's embedded in AI's production process itself. [3] exposes the "digital sweatshops" where workers in Venezuela, Kenya, and the Philippines label data for pennies per task, enabling the AI systems that will later be used to make decisions about their own lives. These workers remain invisible in governance frameworks focused on algorithmic transparency and user privacy.

The concentration of harm documentation in the Global North, combined with the concentration of harm experience in the Global South, creates a discourse that systematically misrecognizes both problems and solutions. When those with the power to study bias aren't those who experience its worst effects, certain forms of injustice become literally unspeakable in the dominant frameworks of AI ethics.

The Silence of the Surveilled

While institutions and experts dominate AI discourse, those subjected to AI's decisions remain largely silent—not because they have nothing to say, but because the structures of discourse provide no channels for their voices. This enforced silence is most acute in education, where students experience AI surveillance, algorithmic grading, and automated discipline while having no meaningful input into these systems' design or deployment.

[17] reveals how students face life-altering accusations based on algorithmic decisions they cannot challenge or understand. The article documents cases where international students faced deportation threats based on false positives from plagiarism detection software. Yet student voices appear in the discourse only as data points or cautionary tales, never as legitimate participants in debates about educational technology.

The surveillance extends beyond academia. [4] details how K-12 students' every digital interaction is monitored by AI systems that flag "concerning" behavior to authorities. The lawsuit reveals that students discussing depression, sexuality, or political views trigger alerts that can lead to disciplinary action or law enforcement involvement. The students had no voice in implementing these systems and little recourse when wrongly flagged.

[16] The Human Side of AI Bias

[3] Detrás del auge de la Inteligencia Artificial hay un ejército de ...

[17] The Unfairness of AI-Flagged Academic Misconduct Investigations in UK ...

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

This systematic silencing extends to other marginalized groups. [19] argues that medical AI fails precisely because it excludes the perspectives of those it claims to help. Patients, particularly those from marginalized communities, become objects of algorithmic analysis rather than subjects with valuable knowledge about their own health needs.

The absence of these voices isn't accidental—it reflects structural power relations that determine who counts as a stakeholder, whose consent matters, and whose expertise is valued. [5] reveals that children increasingly turn to AI chatbots for emotional support and information about sensitive topics. Yet children's perspectives on these tools, their needs, and their vulnerabilities remain absent from policy discussions dominated by parents, educators, and tech companies.

The silencing is maintained through multiple mechanisms. Technical complexity excludes those without specialized knowledge. Institutional channels favor organized stakeholders over diffuse populations. Legal frameworks focus on corporate compliance rather than individual experiences. Most fundamentally, the framing of AI as a technical rather than political issue delegitimizes the expertise of lived experience.

[19] Without Patient Input, AI for Healthcare is Fundamentally Flawed

[5] Los chatbots de IA son la opción predilecta de millones de niños ...

When Transparency Becomes a Burden

One of the most revealing aspects of power in AI discourse is how proposed solutions can become new forms of control. The widespread call for "transparency" and "explainability" in AI systems demonstrates this dynamic, showing how remedies designed by the powerful can create additional burdens for the powerless.

[18] provides a striking example. When gig workers demanded to understand how algorithmic management systems made decisions, companies responded with detailed explanations of every rating factor. But this "transparency" became a new form of stress—workers spent unpaid hours analyzing metrics, optimizing behaviors, and second-guessing every action. The solution addressed regulators' concerns about explainability while creating new forms of worker exploitation.

This pattern repeats across domains. In education, [1] notes that making assessment algorithms "transparent" often means providing teachers with complex statistical outputs they lack time or training to interpret. The transparency requirement is satisfied, but teachers face new administrative burdens without gaining meaningful control over the systems shaping their students' futures.

[18] TMI AI: How too much info from artificial bosses can harm gig workers

[1] Algorithmic Bias in Education

The perverse effects of transparency reveal deeper power dynamics. Those demanding transparency—regulators, advocates, researchers—operate from positions where information translates to agency. They assume that understanding a system enables changing it. But for those subjected to algorithmic decisions, transparency without power merely illuminates their powerlessness. A gig worker who knows exactly why the algorithm cut their pay cannot change the algorithm. A student who understands why they were flagged for cheating cannot challenge the system's authority.

[6] examines this dynamic in the context of EU non-discrimination law. While legal frameworks mandate algorithmic transparency, they don't address the fundamental power imbalance between those who deploy algorithms and those subjected to them. Transparency becomes a procedural requirement that legitimizes continued use of discriminatory systems rather than preventing discrimination.

The focus on transparency also reveals whose concerns shape AI governance. Corporate actors worry about trade secrets and competitive advantage. Regulators seek auditable systems. Researchers want interpretable models. But those experiencing algorithmic decisions often want something simpler: the ability to opt out, to contest decisions, or to access human judgment. These desires find little space in governance frameworks obsessed with making black boxes translucent rather than questioning whether they should exist at all.

Reclaiming Agency in the Age of Algorithms

The analysis of power relations in AI discourse reveals a system designed to maintain existing hierarchies while appearing to address concerns about fairness and justice. From the assumption of inevitability to the governance fixation, from the geographic concentration of harm documentation to the silencing of those surveilled, each aspect of how we talk about AI reinforces the power of those who profit from its expansion while marginalizing those who bear its costs.

Yet understanding these power dynamics also points toward possibilities for resistance and transformation. [11] documents growing pushback against educational AI, with teachers, parents, and students questioning not just how AI is implemented but whether it should be used at all. [7] shows that sustained activism can achieve concrete policy victories, rejecting surveillance technologies rather than merely regulating them.

The path forward requires fundamental shifts in how we approach AI discourse. Instead of asking how to make AI fair, we must ask

[6] Lutter contre les discriminations produites par les algorithmes et l'I ...

[11] Report: The risks of AI in schools outweigh the benefits : NPR

[7] New York bans facial recognition in schools after report finds risks ...

who decides what fairness means. Instead of seeking transparency, we must demand agency. Instead of mitigating bias, we must question the authority of algorithmic judgment itself. Most critically, we must center the voices of those most affected by AI systems—not as data points or user feedback, but as legitimate experts on their own lives with the power to refuse, reshape, or reject the technologies imposed upon them.

The discourse on AI and social aspects reveals that the greatest threat isn't biased algorithms or opaque systems—it's the concentration of power to define problems, impose solutions, and shape futures without meaningful input from those whose lives hang in the balance. Recognizing these power relations is the first step toward challenging them. The question isn't whether AI will transform society, but whether those transformations will be dictated by the powerful or determined democratically by all those affected.

As [12] argues, existing legal frameworks are inadequate to address algorithmic discrimination because they assume human decision-makers and individual cases rather than systematic, automated bias. But perhaps the deeper inadequacy lies in trying to regulate our way to justice while leaving fundamental power relations intact. Until those subjected to algorithmic decisions have real power to shape, refuse, or dismantle these systems, all the transparency, explainability, and governance frameworks in the world will merely legitimate continued oppression under the guise of responsible innovation.

The critical task ahead isn't to make AI more ethical or transparent—it's to democratize the power to decide whether, when, and how algorithmic systems shape human lives. That requires not just better governance but fundamental shifts in who governs, whose knowledge counts, and whose consent matters in the age of algorithms.

[12] Resetting Antidiscrimination Law in the Age of AI

References

1. Algorithmic Bias in Education
2. Assessing the AI Moment from a Global South Perspective
3. Detrás del auge de la Inteligencia Artificial hay un ejército de ...
4. Lawrence school district sued in federal court for use of AI-powered
...
5. Los chatbots de IA son la opción predilecta de millones de niños ...
6. Lutter contre les discriminations produites par les algorithmes et l'I
...

7. New York bans facial recognition in schools after report finds risks
...
8. PDF Are Algorithms Biased in Education? Exploring Racial Bias in Predicting ...
9. PDF Rapport algorithmes, systèmes d IA et services publics : quels droits ...
10. Predictive policing algorithms are racist. They need to be dismantled.
11. Report: The risks of AI in schools outweigh the benefits : NPR
12. Resetting Antidiscrimination Law in the Age of AI
13. South Korea's 'world-first' AI laws face pushback amid bid to become leading tech power
14. Students have been called to the office for AI surveillance false alarms
15. The cultural cost of AI in Africa's education systems - UNESCO
16. The Human Side of AI Bias
17. The Unfairness of AI-Flagged Academic Misconduct Investigations in UK ...
18. TMI AI: How too much info from artificial bosses can harm gig workers
19. Without Patient Input, AI for Healthcare is Fundamentally Flawed