AI and Social Aspects

Weekly Analysis — https://ainews.social

The Architecture of Silence: Power and Discourse in AI's Social Order

The conversation about artificial intelligence unfolds in predictable patterns. Institutions speak to institutions. Experts debate with experts. Meanwhile, those whose lives are upended by algorithmic systems—welfare recipients denied benefits, students flagged by surveillance systems, data laborers traumatized by content moderation—remain conspicuously absent from the discourse that determines their fate. This silence is not accidental. It is the architecture of power.

The evidence is stark: across 1,486 articles examining AI's social aspects, ethical failures dominate the narrative at 41.3%, yet solutions remain elusive. As [18] documents, algorithmic systems are causing severe harm to vulnerable populations across welfare systems, education, and employment. But who shapes the response? Who decides what counts as a problem worth solving? The answer reveals a discourse controlled by those least affected by AI's harms and most invested in its continued expansion.

[18] This Algorithm Could Ruin Your Life - WIRED

This essay interrogates the power relations embedded in AI discourse, exposing how institutional voices drown out those of the harmed, how technical frames obscure political choices, and how the geography of AI development perpetuates colonial extraction. By examining who speaks and who remains silent, who frames problems and who suffers consequences, we can begin to understand not just what AI does, but what the discourse about AI accomplishes—namely, the maintenance of existing power structures under the guise of innovation.

The Institutional Echo Chamber

The dominance of institutional perspectives in AI discourse is overwhelming. Scan any collection of AI governance discussions and you'll find a predictable cast: policymakers crafting frameworks, legal professionals navigating compliance, corporate officers managing risk.

These voices shape what counts as legitimate concern and reasonable response. [Colorado releases new AI Policy framework aimed revising the state's 2024 law] exemplifies this pattern—institutions speaking to institutions about institutional concerns, with affected communities nowhere in sight.

This institutional capture extends across domains. In education, administrators and policymakers debate AI integration while students subjected to surveillance systems have no platform. [16] reveals the human cost of these systems—students wrongly flagged, pulled from class, subjected to disciplinary procedures based on algorithmic errors. Yet when solutions are proposed, they emerge from boardrooms and committee meetings, not from those experiencing the technology's sharp edge.

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

The pattern repeats globally. [9] shows European institutions determining the legal architecture for AI development, while the Global South workers actually training these systems through grueling data labeling remain voiceless. The discourse privileges those who build frameworks over those who bear consequences.

[9] European Parliament Study Recommends Statutory Licensing as the Optimal Copyright Framework for AI Training

What makes this institutional dominance particularly insidious is its self-reinforcing nature. Institutions define the problems worth addressing—typically framed as governance challenges requiring technical or regulatory solutions. They set the terms of debate, establish what counts as evidence, and determine whose expertise matters. Missing from these carefully managed conversations are the messy realities of algorithmic harm: the welfare recipient whose benefits were terminated by an opaque algorithm, the student whose educational trajectory was altered by a biased prediction model, the content moderator developing PTSD from training AI systems.

Victims Without Voices

The most striking feature of AI discourse is the systematic absence of those most affected by algorithmic systems. [11] provides a rare exception, documenting how algorithmic fraud detection systems across Europe systematically target and harm vulnerable welfare recipients. The investigation required accessing code and data to reveal what those affected already knew: the systems were discriminatory by design, not by accident.

[11] Inside the Suspicion Machine - WIRED

This pattern—where investigative journalism must excavate what affected communities experience daily—reveals the discourse's fundamental flaw. Those subjected to algorithmic decision-making rarely have platforms to share their experiences. When [15] exposed how

[15] Scoring of welfare beneficiaries: the indecency of CAF's algorithm now undeniable

France's welfare algorithm systematically discriminated against single mothers and residents of disadvantaged neighborhoods, it took years of advocacy and legal action to make visible what thousands had been experiencing.

In educational contexts, the silencing is equally pronounced. [5] demonstrates that predictive models systematically underserve marginalized students, leading to fewer resources and support for those who need them most. Yet education policy discussions remain dominated by administrators and vendors promoting AI solutions, not by the students and families experiencing algorithmic discrimination.

[5] Are algorithms biased in education? Exploring racial bias in predicting ...

The absence extends to labor. While discourse celebrates AI's transformative potential, [1] reveals the hidden reality: armies of workers in Kenya, India, and other Global South countries performing the traumatic work of content moderation and data labeling that makes AI possible. These workers report developing PTSD from exposure to violent content, earning wages below subsistence levels, yet their experiences rarely penetrate mainstream AI discourse.

[1] 'A.I. Is African Intelligence': Workers Who Train A.I. Fight Back

What does appear are sanitized accounts, statistical aggregations that transform human suffering into manageable metrics. [19] represents a laudable attempt to capture public sentiment, yet even this massive study cannot fully capture the experiences of those most harmed by AI systems—those often too marginalized to participate in such research.

[19] What 81000 people want from AI

The Geography of Power

The global dimension of AI discourse reveals stark power asymmetries. Western institutions—particularly in the US and Europe—dominate both AI development and the discourse about its governance, while the Global South bears disproportionate costs. [6] articulates this clearly: AI development follows colonial patterns of resource extraction, with data and labor flowing from South to North while benefits concentrate in already-powerful centers.

[6] Artificial intelligence, digital colonialism, and the implications for Africa's future development

This geographic inequality manifests in concrete ways. Africa represents only 1% of global data center capacity and produces just 0.77% of global AI publications, yet African data and labor are essential to AI development. [17] exposes how Western-developed educational AI systems impose foreign pedagogical approaches and erase local knowledge systems, a form of algorithmic colonialism in the classroom.

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

The regulatory landscape reflects these power imbalances. [14] reveals how African nations, lacking resources to develop comprehen-

[14] Pourquoi la protection des données est devenue l'outil de politique d'IA par défaut en Afrique

sive AI policies, default to data protection frameworks modeled on European GDPR—regulations designed for different contexts and capabilities. This regulatory dependence perpetuates technological dependence.

The human cost of this geographic inequality is devastating. [New Report Series on Human Labour Powering AI] documents how content moderators and data laborers in the Global South suffer severe psychological trauma from their work training AI systems, yet their experiences remain invisible in Silicon Valley boardrooms where AI strategy is determined. The discourse celebrates AI's potential while hiding its human infrastructure.

Even attempts to address these inequalities often reinforce them. Development initiatives frame AI as a solution to Global South challenges without questioning who controls the technology or benefits from its deployment. The discourse positions the Global South as recipients of AI beneficence rather than as equal participants in shaping AI's development and governance.

Laundering Accountability Through Abstraction

One of power's most effective strategies in AI discourse is the transformation of concrete harms into abstract problems requiring technical solutions. When algorithmic systems deny welfare benefits, exclude job applicants, or flag students for punishment, these specific injustices become "bias"—a technical problem requiring technical fixes. [3] challenges this framing, arguing that algorithmic bias is fundamentally a social and political problem, not merely a technical one.

[3] Algorithmic bias and discrimination through digitalization in education: a socio-technical view

This abstraction serves multiple functions. First, it diffuses responsibility. When bias is framed as an emergent property of complex systems rather than the predictable result of specific choices, no one is accountable. [10] documents systematic administrative errors in public benefits systems, yet these errors are treated as unfortunate technical glitches rather than as violence enacted through administrative means.

[10] How Do Algorithmic Decision-Making Systems Used in Public Benefits ...

Second, technical framing privileges technical expertise. If bias is a problem of data representation and model optimization, then data scientists and engineers become the natural authorities on solutions. Missing from this conversation are social workers who understand poverty's complexities, educators who grasp learning's nuances, or communities who experience discrimination's daily reality. [8] attempts to bridge this gap but remains largely within technical discourse.

[8] Debiasing Education Algorithms | International Journal of Artificial ...

Third, abstraction enables the continuation of harmful practices under the guise of improvement. [7] exemplifies this: rather than preventing discriminatory systems, regulation focuses on transparency—making visible the discrimination that continues to occur. Power maintains itself by appearing to reform.

[7] Colorado Moves to Replace AI Law's Bias Audit Requirements With Transparency Framework: 5 Action Steps for Employers

The consequences of this abstraction are material. While experts debate fairness metrics and transparency frameworks, people lose benefits, students face unjust punishment, and workers endure exploitation. The discourse's abstract nature isn't a bug—it's a feature that allows harmful systems to continue operating while appearing to address concerns.

The Convenient Invisibility of Labor

Perhaps nowhere is power's hand more visible than in the systematic erasure of AI's human labor infrastructure from mainstream discourse. AI systems are presented as autonomous, almost magical technologies, obscuring the massive human workforce required for their development and maintenance. This invisibility is politically convenient—it allows the AI industry to maintain narratives of efficiency and automation while relying on exploitative labor practices.

[L'IA et la fin du temps] provides a philosophical framework for understanding this erasure, examining how AI transforms labor and value while hiding its own dependence on human work. The article's analysis of "temporal regimes" reveals how AI discourse constructs a future-oriented narrative that renders present exploitation invisible.

The exploitation is staggering in scope and severity. Content moderators across the Global South spend their days viewing the worst of human behavior—violence, abuse, extremism—to train AI systems to recognize harmful content. These workers, essential to making AI systems "safe," develop severe psychological trauma. Yet their existence barely registers in discourse about AI ethics and safety, which focuses on abstract principles rather than concrete working conditions.

[2] addresses one dimension of AI's labor politics—surveillance of workers by AI systems—but even this critical work maintains the separation between those who build AI and those subjected to it. Missing is recognition that these categories overlap: the same workers who label data for AI development are surveilled by AI systems in their workplaces.

[2] A policy primer and roadmap on AI worker surveillance and ... - Springer

This invisibility extends to educational contexts. While institutions debate AI's pedagogical potential, they rely on adjunct instructors

and teaching assistants to manage the actual work of education—labor made more precarious by AI's introduction. The discourse celebrates personalized learning while ignoring how such systems depend on educators' invisible labor to function.

The convenience of this invisibility for power is obvious. By hiding AI's labor infrastructure, the industry avoids accountability for working conditions, wages, and psychological harm. It maintains the fiction of technological autonomy while depending on human exploitation. And it shifts discourse toward future promises rather than present realities.

Reframing Solutions, Maintaining Power

The most insidious aspect of power in AI discourse is how proposed solutions often reinforce existing hierarchies rather than challenging them. Across domains, we see a consistent pattern: reforms that appear responsive to criticism while maintaining fundamental power structures intact.

Consider transparency, the most common prescription for algorithmic accountability. [13] extensively documents the need for transparency in public sector AI use. Yet transparency without power to act on revealed information is merely documentation of ongoing harm. Showing people how they're being discriminated against doesn't stop the discrimination.

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

Similarly, calls for "human-in-the-loop" systems maintain the fiction that adding human oversight solves algorithmic problems. But when humans lack power to meaningfully intervene—when case workers must process hundreds of cases using algorithmic recommendations, when teachers are evaluated by systems they can't challenge—human involvement becomes a fig leaf for automated injustice.

The focus on technical debiasing reveals the same dynamic. [4] comprehensively reviews bias in educational AI, and [12] extends this critique to generative AI. Yet technical debiasing assumes the problem is imperfect implementation rather than questioning whether these systems should exist at all. Power maintains itself by constraining discourse to how to fix harmful systems rather than whether to use them.

[4] Algorithmic Bias in Education - Academia.edu
[12] Major Concerns of Generative AI in Education: A Critique

The emphasis on stakeholder engagement and multi-stakeholder governance sounds democratic but often reinforces existing inequalities. When stakeholders meet, who has a seat at the table? Who has resources to participate meaningfully? Who understands the technical

language that dominates discussion? [Detroiters cautious about government use of AI, U-M survey ...] reveals public skepticism about AI use, yet public input rarely shapes deployment decisions.

Even critical discourse can maintain power structures by focusing on reform rather than transformation. Discussions of ethical AI, responsible AI, and trustworthy AI assume AI's continued development and deployment—the question becomes how to do it better, not whether to do it at all. Power wins by setting the terms of debate.

Conclusion: Toward a Politics of Refusal

The architecture of AI discourse reveals power operating through multiple mechanisms: institutional capture, systematic silencing, geographic inequality, convenient abstraction, labor invisibility, and reformist solutions that maintain the status quo. Understanding these mechanisms is essential but insufficient. The question becomes: how do we build counter-power?

The answer cannot come from within existing discourse structures. As [Are humans being left behind in the artificial intelligence push?] suggests, current AI development prioritizes technological capability over human flourishing. Real change requires amplifying voices systematically excluded from AI discourse—welfare recipients, surveilled students, exploited workers, Global South communities whose data and labor power AI while receiving few benefits.

It requires rejecting the technical framing of political problems. When [3] argues for socio-technical approaches, it opens space for recognizing that algorithmic harm isn't a bug to be fixed but a feature of systems designed to automate and scale existing inequalities. The solution isn't better algorithms but different social arrangements.

[3] Algorithmic bias and discrimination through digitalization in education: a socio-technical view

Most radically, it may require a politics of refusal—rejecting the premise that all domains of human life should be subject to algorithmic mediation. When every proposed solution maintains AI's centrality while tweaking its implementation, power has already won the essential battle. The discourse needs voices willing to say: not here, not this, not now. The students refusing surveillance, the workers organizing against algorithmic management, the communities rejecting algorithmic governance—these are the voices that might crack open the closed loop of institutional AI discourse.

The 41.3% of articles documenting ethical failures in AI systems tell us something important: the current approach isn't working. But they tell us something more important through their silences—through who

isn't speaking, whose solutions aren't considered, whose refusal isn't heard. Until AI discourse centers those most harmed by algorithmic systems, until it takes seriously the possibility that some domains should remain beyond algorithmic reach, it will continue to serve power rather than challenge it.

The path forward requires more than better governance frameworks or technical fixes. It requires fundamentally reimagining who has the right to shape our algorithmic futures—and who has the right to refuse them altogether. Only by breaking the institutional monopoly on AI discourse, by amplifying systematically silenced voices, and by taking seriously the politics of refusal can we begin to build a more just relationship with these powerful technologies. The alternative is what we have now: a discourse that documents harm while enabling its continuation, that speaks of ethics while perpetuating exploitation, that promises transformation while maintaining existing hierarchies of power.

In the end, the question isn't how to fix AI but how to fix the power relations that shape its development and deployment. That work begins by listening to those the current discourse systematically silences—and by taking seriously their demands not just for inclusion but for the right to refuse algorithmic mediation altogether. Only then can we move beyond an AI discourse that serves power toward one that serves people.

References

1. 'A.I. Is African Intelligence': Workers Who Train A.I. Fight Back

2. A policy primer and roadmap on AI worker surveillance and ... - Springer

3. Algorithmic bias and discrimination through digitalization in education: a socio-technical view

4. Algorithmic Bias in Education - Academia.edu

5. Are algorithms biased in education? Exploring racial bias in predicting ...

6. Artificial intelligence, digital colonialism, and the implications for Africa's future development

7. Colorado Moves to Replace AI Law's Bias Audit Requirements With Transparency Framework: 5 Action Steps for Employers

8. Debiasing Education Algorithms | International Journal of Artificial ...

9. European Parliament Study Recommends Statutory Licensing as the Optimal Copyright Framework for AI Training

10. How Do Algorithmic Decision-Making Systems Used in Public Benefits ...

11. Inside the Suspicion Machine - WIRED

12. Major Concerns of Generative AI in Education: A Critique

13. PDF Rapport algorithmes, systèmes d IA et services publics : quels droits ...

14. Pourquoi la protection des données est devenue l'outil de politique d'IA par défaut en Afrique

15. Scoring of welfare beneficiaries: the indecency of CAF's algorithm now undeniable

16. Students have been called to the office for AI surveillance false alarms

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

18. This Algorithm Could Ruin Your Life - WIRED

19. What 81000 people want from AI