International Journal of Science, Technology and Society
2026, Vol. 14, No. 2, pp. 88–101
https://doi.org/10.11648/j.ijsts.20261402.15

Research Article

The Ghost Workforce of Generative AI: Reclaiming Ethical Visibility in AI Labour

Achi Iseko*

Independent Researcher, Surrey, United Kingdom

Abstract

Generative AI systems are widely framed as autonomous technologies, yet they depend on extensive human labour data annotation, content moderation and reinforcement learning—which remains largely invisible within dominant narratives of automation and AI ethics. Drawing on 80 semi-structured interviews across eight countries in Sub-Saharan Africa and the Arab Gulf, this study conceptualises generative AI as a labour-intensive communicative infrastructure and examines how invisibility is systematically produced in AI supply chains. This article introduces the Ethical Visibility Framework (EVF), which identifies three interlocking forms of invisibility: material invisibility, enacted through outsourcing, contractual opacity and wage suppression; epistemic invisibility, whereby workers' situated knowledge is extracted while their status as legitimate contributors is denied; and narrative invisibility, sustained through automation discourse that erases human labour from accounts of AI performance. These dimensions reinforce one another, enabling economic extraction, expertise appropriation and responsibility deflection. By centring AI workers as epistemic agents rather than peripheral service providers, the study extends scholarship on digital labour, data justice and critical AI ethics, concluding that labour visibility is a foundational ethical requirement for legitimate AI development, not an ancillary concern, with implications for governance frameworks that treat AI ethics as technical optimisation rather than socio-technical justice.

Keywords

Generative AI, AI Labour, Ethical Visibility, Epistemic Justice, Content Moderation, Platform Work

1. Introduction: Rethinking Ethics in the Age of Generative AI

Artificial intelligence systems do not emerge in isolation; they are embedded within socio-technical assemblages shaped by power, labour and cultural narratives. Critical scholarship in Science and Technology Studies (STS) has long emphasised that technologies are not neutral tools but artefacts which reproduce social hierarchies and imaginaries (Suchman [52]; Bowker & Star [7]). However, much of contemporary AI ethics discourse remains narrowly focused on technical metrics of fairness, accuracy or transparency, while overlooking the hidden infrastructures of human labour that make AI possible. This paper intervenes in these debates by foregrounding the ethical invisibility of labour in AI systems as a societal problem, not merely a technical or managerial concern.

*Correspondence: Achi Iseko (a.iseko@yahoo.com)

Received: 16 March 2026; **Accepted:** 1 April 2026; **Published:** 28 April 2026

Since Turing [55] first asked whether machines can think, AI research has centred on questions of machine intelligence while removing the human labour sustaining these systems. Mainstream textbooks present AI as autonomous agents, detached from the material and social conditions of their production [44]. Yet, as Crawford [14] argues, artificial intelligence is always embedded in extractive infrastructures—ecological, economic and human. This study examines the "ghost workforce" underpinning generative AI: data annotators, content moderators and reinforcement-learning contractors whose cognitive and emotional labour trains and aligns large language models, yet who remain structurally invisible within dominant narratives of automation and innovation [2, 29, 48].

The invisibility of human contributions to AI is not incidental but structurally produced. Workers are outsourced, geographically displaced and discursively erased from narratives of innovation, reinforcing colonial patterns of extraction and epistemic silencing (Birhane [4]; Milan & Treré [37]). While digital labour studies have illuminated aspects of these dynamics (Gray & Suri [23]; Irani [28]; Roberts [43]), a systematic conceptual lens remains lacking to theorise how invisibility is constructed and sustained across material, epistemic and narrative dimensions of AI production.

This article introduces the Ethical Visibility Framework (EVF) as a response to that gap. Situated at the intersection of AI ethics, labour studies and STS traditions, EVF provides a structured lens to examine the processes through which human labour in AI is rendered invisible. It conceptualises invisibility not as a passive absence but as an active condition of technological production shaped by platform capitalism, epistemic hierarchies and automation discourse. By reframing visibility as an ethical precondition for justice, EVF advances the argument that responsible AI must address not only algorithmic outcomes but also the recognition and participation of those who make AI systems possible.

Methodologically, the paper draws on 80 semi-structured interviews with data annotators, content moderators, in-house AI ethics specialists and policy consultants across eight countries in Sub-Saharan Africa and the Arab Gulf (see Table 1). This multi-sited design grounds EVF's conceptual development while illuminating how invisibility is enacted and resisted across diverse regional contexts [20, 61].

The theoretical approach draws on critical traditions that foreground power and epistemology. Haraway [25] "situated knowledges" highlight the partial, embodied nature of all knowledge claims, while Birhane [5] relational ethics emphasises interdependence and context. Infrastructural perspectives show how classification systems render certain forms of work invisible [7] and decolonial critiques stress the extractive logic of computational sciences [6]. These perspectives inform EVF's central move: positioning annotators, moderators and RLHF contractors not as peripheral service providers but as epistemic agents whose situated expertise is systematically harvested while their status as knowledge producers is denied.

The contribution of this paper is twofold: it introduces EVF as a novel framework for theorising invisibility in AI labour through three interconnected dimensions—material, epistemic and narrative—demonstrating how they constitute a mutually reinforcing system and then applies this framework to empirical cases, revealing patterns of structural precarity, expertise extraction and discursive erasure across AI supply chains.

Section 2 reviews existing scholarship on AI labour, organising literature around EVF's three dimensions. Section 3 outlines the methodological design. Section 4 presents empirical findings structured around EVF's dimensions. Section 5 discusses theoretical contributions and policy implications. Section 6 concludes by reflecting on the ethical urgency of making AI labour visible and proposes directions for future research and governance reform.

2. Literature Review: Theorising Invisibility in AI Labour

The rapid expansion of artificial intelligence has been accompanied by a paradox: while AI is framed as autonomous intelligence, it remains profoundly dependent on human labour. This literature review organises existing scholarship around three dimensions through which invisibility is produced and maintained: material structures of labour precarity, epistemic hierarchies that devalue worker knowledge and narrative framings that discursively erase human contributions. Together, these dimensions form the foundation for the Ethical Visibility Framework (EVF).

2.1. Material Invisibility: Political Economy of AI Labour

The invisibility of AI labour is first and foremost a structural condition rooted in the global political economy. Platform capitalism has enabled unprecedented geographic dispersal and organisational fragmentation of work, creating supply chains that obscure the human foundations of AI systems [49]. What Miceli and Posada [34] term the "data-production dispositive" systematically organises, disciplines and extracts value from data workers while rendering their contributions invisible. This is not incidental oversight but strategic governance: by placing workers in different legal jurisdictions, time zones and employment relationships, AI companies distance accountability and diffuse responsibility.

Recent scholarship has begun mapping the full spectrum of human labour in AI systems. Catanzariti et al. [10] and Lampinen et al. [31] document how "AI work" extends far beyond software engineering to include data collection, annotation, content moderation and quality assurance—roles typically outsourced to Global South workers [22]. Gray and Suri [23] ethnography of "ghost work" demonstrates how invisible labour sustains digital platforms yet is deprived of recognition, stability or dignity. Tubaro, Casilli and Coville [54] similarly describe a global workforce whose contributions are neither acknowledged in corporate discourses nor protected under labour law.

This labour invisibility serves economic functions, functioning as what Pasquinelli and Joler [42] term knowledge extractivism — the harvesting of human cognition at scale without attribution or compensation. As Sambasivan et al. [45] demonstrate, organisational emphasis on "model work" over "data work" enables hierarchical compensation structures where data workers earn a fraction of ML engineers' salaries despite comparable cognitive demands. The result is what they term "data cascades": compounding errors stemming from undervalued, under-resourced data labour which systematically degrades AI system quality. However, these failures are attributed to "technical issues" rather than labour conditions, further obscuring workers' contributions [9].

The political economy of AI labour mirrors earlier industrial capitalism, where supply chains concealed exploitative conditions (Couldry & Mejias [13]). The difference lies in AI labour's immateriality: it is cognitive, affective and cultural, yet commodified through platforms that fragment tasks into "micro-work" and suppress worker identities. Material invisibility is thus not passive absence but active production [60]. Firms present AI as autonomous innovation, legitimising monopolistic rents while deflecting demands for labour protections [1]. Recent work on workplace democracy in data work [19] suggests that making work visible is a precondition for workers exercising collective power. Without structural transparency in AI supply chains, ethical governance remains impossible [41, 50].

2.2. Epistemic Invisibility: Knowledge Hierarchies in AI Development

Beyond economic marginalisation, AI workers face systematic exclusion from epistemic communities which define expertise, innovation and legitimate knowledge in AI development. This operates through what Fricker [21] terms "epistemic injustice": the systematic devaluation of knowledge claims based on speaker identity rather than content validity. In AI systems, workers' experiential expertise—judgments about language nuance, cultural context and ethical boundaries—is treated as noise rather than legitimate contribution, while engineers and researchers are positioned as sole knowledge producers.

This epistemic hierarchy has material consequences. Sambasivan and Veeraraghavan [46] document "deskilling" in AI development: the systematic devaluation of domain expertise that data workers bring from agriculture, healthcare, education and other fields. When annotators with agricultural knowledge label crop disease images or nurses classify medical scans, their domain expertise is harvested while their status as knowledge producers is denied. The result is what the authors call "expertise laundering"—situated knowledge extracted from workers and repackaged as algorithmic intelligence attributed solely to model developers.

The politics of invisibility in AI data work extend beyond individual projects to shape whose knowledge counts in defining AI ethics itself. As Miceli, Posada and Yang [35] argue, dominant framings of "algorithmic bias" obscure underlying power relations by focusing on technical model properties rather than the political economies and labour conditions that produce datasets. The shift from bias to power reframes epistemic questions as fundamentally about recognition and participation.

Postcolonial and decolonial scholarship deepens this analysis by situating epistemic erasure within longer histories of colonial knowledge extraction. Mbembe [33] and Ndlovu-Gatsheni [39] emphasise how epistemic domination shapes knowledge production: Northern institutions monopolise the authority to define "ethical AI" thereby marginalising alternative epistemologies, including African philosophies of relationality or Latin American traditions of liberation ethics [36]. Birhane [4] argues that relational, context-sensitive ethics provides critical counterpoints to the abstract universalism of dominant AI frameworks. Yet workers in the Global South—who grapple daily with cultural nuance and contextual judgment—are systematically excluded from governance processes that claim to address "fairness" and "inclusion" [57].

2.3. Narrative Invisibility: Metaphors and Discursive Erasure

Invisibility is not only enacted through structures and knowledge systems but also reproduced through language, metaphors and dominant narratives about AI. Terms like "machine learning" and "neural networks" linguistically erase human teachers and cognitive architects. Corporate communications describe "self-learning AI" without mentioning annotation workforces [15, 40].

The automation mythology surrounding AI performs significant discursive work. By framing AI as autonomous agent rather than socio-technical assemblage, dominant narratives obscure the human labour sustaining these systems. As Star and Strauss [51] demonstrate in their analysis of infrastructural invisibility, labour supporting technological "front ends" is systematically overlooked—most visible in breakdown, when errors reveal human fingerprints behind supposedly autonomous systems. In AI, this means workers become perceptible only through failure: when content moderation misses harmful material, when annotation errors produce biased outputs and when system breakdowns expose supply chains [2, 26].

This discursive erasure extends to responsibility frameworks. When AI systems produce harmful outputs, discourse focuses on "algorithmic bias" or "model errors" rather than the labour conditions of traumatised moderators who reviewed training data or underpaid annotators working without adequate cultural context. Greene, Hoffmann and Stark [24] argue that many responsible AI initiatives function as "ethics-washing"—rhetorical commitments that obscure structural inequities. By focusing on model fairness or technical explainability, these frameworks sidestep precarious labour conditions while claiming ethical legitimacy [12].

The proliferation of responsible AI frameworks—from the EU AI Act to OECD Principles—emphasises transparency, accountability and fairness [59]. Yet these frameworks often neglect material labour. Ethical guidelines reference "human oversight" without acknowledging the global workforce performing such oversight under precarious conditions. This disconnect reveals a deeper problem: governance efforts that treat AI ethics as technical optimisation rather than socio-technical justice will inevitably reproduce the invisibilities they claim to address.

Making AI labour narratively visible requires fundamentally reframing how we talk about AI systems—not as autonomous agents but as assemblages of distributed human and machine labour, with corresponding obligations to recognise and protect human contributors.

2.4. Toward Ethical Visibility: Synthesising a Framework

AI labour's invisibility is neither inevitable nor benign. It is structurally produced through platform capitalism and supply chain fragmentation, epistemically entrenched through expertise hierarchies and knowledge extraction, and discursively reinforced through automation narratives and metaphors.

This study introduces the Ethical Visibility Framework (EVF) to address that gap. Drawing from Science and Technology Studies (Suchman [52]; Bowker & Star [7]), postcolonial theory (Birhane [4]; Ndlovu-Gatsheni [39]), intersectionality [16] and political economy (Srnicek [49]; Couldry & Mejias [13]), EVF synthesises insights across these traditions through the unifying problematic of visibility. Haraway [25] "situated knowledges" inform EVF's epistemological stance, while Birhane [5] relational ethics emphasises the interdependence and context-specificity that dominant AI frameworks abstract away.

EVF articulates three interlocking dimensions through which invisibility operates:

1) *Material invisibility* recognises how political economy, platform intermediation and supply chain structures render workers physically and contractually distant from AI systems they produce, enabling wage suppression and rights denial.
2) *Epistemic invisibility* examines how expertise hierarchies and knowledge extraction systems harvest workers situated knowledge while denying their status as legitimate knowledge producers in AI governance and design.
3) *Narrative invisibility* analyses how automation discourse, corporate communications and governance frameworks discursively erase human contributions, deflecting ethical responsibility from labour conditions to abstract technical properties.

Critically, EVF demonstrates that these dimensions constitute a mutually reinforcing system. Material precarity is justified through epistemic exclusion (workers framed as low-skill, substitutable), which is naturalised through narrative (automation discourse positions human work as temporary transitional labour). This cycle makes invisibility self-sustaining—interrupting it requires coordinated interventions across all three dimensions.

Critically, EVF demonstrates that these dimensions constitute a mutually reinforcing system. Material precarity is justified through epistemic exclusion, which is naturalised through narrative — a self-sustaining cycle examined empirically in Section 4.4. The framework offers both an analytical lens to examine how invisibility is produced and a normative foundation for justice-oriented AI governance [15, 35].

3. Methodology

This study employed a qualitative, multi-sited design to examine hidden labour practices underpinning generative AI. Following critical interpretivist traditions that privilege lived experience and situated knowledge (Haraway [25]; Birhane [5]), the research centred workers' voices as legitimate epistemic contributions. The methodology was deliberately designed to bridge the Global North-South divide by conducting fieldwork across eight countries in Sub-Saharan Africa and the Arab Gulf.

3.1. Research Design and Participant Recruitment

The study draws on 80 semi-structured interviews conducted between 2021 and 2024 with workers directly involved in generative AI workflows: data annotation, content moderation, reinforcement learning from human feedback (RLHF), in-house AI ethics and policy consultation. Participants were distributed across eight countries—five in Sub-Saharan Africa and three in the Arab Gulf—reflecting the global division of labour in AI wherein data-intensive work is disproportionately outsourced to the Global South (Irani [28]; Roberts [43]).

Table 1. Participant Overview.

Role Category	N	Geographic Distribution	Primary Tasks	Platform/Company Knowledge
Content Moderators	28	Kenya (9), Nigeria (7), Tanzania (4), Ghana (3), UAE (3), Qatar (2)	Image/video review, harmful content flagging, CSAM detection, hate speech moderation	18 knew specific platform (e.g., Meta, TikTok); 10 worked through intermediaries without end-client knowledge
Data Annotators (RLHF)	24	Kenya (8), South Africa (6), Nigeria (4), Bahrain (3), Qatar (3)	Conversational preference ranking, response quality assessment, linguistic alignment	14 knew they trained LLMs (specifically ChatGPT, Claude, or Bard); 10 understood work as generic "AI training"
Data Annotators (General)	16	South Africa (5), Ghana (4), Kenya (3), UAE (2), Nigeria (2)	Image classification, entity recognition, bounding boxes, text labelling	4 knew specific use case; 12 had no knowledge of end application
In-house AI Ethics Specialists	8	South Africa (3), Kenya (2), Nigeria (2), UAE (1)	Fairness auditing, policy development, internal advocacy, bias testing	Full knowledge of models and deployment contexts
Policy Consultants	4	South Africa (2), Kenya (1), Bahrain (1)	Regional AI governance, regulatory frameworks, civil society liaison	Full knowledge of policy landscape and stakeholder ecosystems
TOTAL	80	Africa: 65 (Kenya=23, South Africa=16, Nigeria=15, Ghana=7, Tanzania=4) Gulf: 15 (UAE=6, Qatar=5, Bahrain=4)		

This geographic distribution captures regions central to AI labour yet marginal to AI ethics discourse. The Gulf states represent an under-examined site positioned between Global North wealth and Global South labour conditions, where workers hold higher formal qualifications but face similar precarity.

Participants were recruited through three channels: direct outreach via LinkedIn (n=38), snowball sampling through worker collectives (n=29) and gatekeeper access through labour organizations (n=13). Of 127 workers contacted, 47 declined citing non-disclosure agreements, surveillance fears and time constraints. To mitigate concerns, interviews used encrypted channels, participant-controlled recording and flexible scheduling.

Participants were not financially compensated to avoid coercive dynamics given workers' economic vulnerability and minimising surveillance risks from payment trails and worker preference for research outputs over monetary compensation. This approach followed ESRC ethical guidelines and received approval as voluntary participation research, though we acknowledge potential selection bias toward economically stable workers.

3.2. Data Collection and Analysis

Interviews lasted 60-90 minutes following a semi-structured guide exploring work routines, compensation perceptions, affective responses, employer relationships and concerns about AI's societal impact. Interviews were conducted in participants' preferred languages, with bilingual research assistants providing interpretation. All interviews occurred remotely via Zoom between January 2022 and March 2024, enabling access to geographically dispersed participants while precluding ethnographic observation.

Interviews were supplemented by analysis of 45 documents including platform terms of service, corporate transparency reports, dataset documentation, labour advocacy reports and policy frameworks from three regulatory contexts (EU AI Act drafts, Kenya Data Protection Act and South Africa AI policy proposals). This enabled triangulation and revealed discrepancies between corporate self-representation and worker experiences.

Transcripts were coded using NVivo 12 following iterative thematic analysis [8]. Open coding identified 47 first-order codes (time pressure, toxic exposure, wage uncertainty, NDA constraints). Axial coding clustered these into 12 second-order themes (material precarity, informational opacity, expertise devaluation, narrative erasure). Selective coding organised themes into three dimensions—material, epistemic, narrative—which became the Ethical Visibility Framework's architecture. Critically, EVF emerged from rather than preceding the data, ensuring the framework reflected lived experiences rather than imposing predetermined categories [47].

To enhance rigor, 25% of transcripts (n=20) were dual-coded, achieving inter-coder agreement of Cohen's $\kappa = 0.81$

for first-order codes and κ = 0.78 for second-order themes. Disagreements were resolved through deliberative discussion. Ten participants engaged in member-checking, reviewing thematic summaries and suggesting refinements that strengthened interpretation.

3.3. Ethics, Positionality and Study Limitations

The study implemented robust participant protection: all participants chose pseudonyms; employer names were omitted or generalised; and geographic details were aggregated where requested. For content moderators describing graphic material exposure, interviews included breaks and exit options, with information provided about accessible psychological support services. Data management followed stringent protocols: encrypted storage, no cloud services, restricted access and planned destruction five years post-publication [32].

As an independent researcher with academic training in information systems and prior industry experience in AI product development, I acknowledge positional privileges shaping access and interpretation. My industry background afforded credibility with in-house specialists and consultants while potentially creating distance from platform workers. I approached this research committed to documenting and amplifying worker perspectives rather than extracting knowledge for purely academic purposes, maintaining reflexive analytical memos documenting assumptions and discomforts throughout fieldwork [11].

Study limitations include geographic concentration in Anglophone Sub-Saharan Africa and select Gulf states, potentially underrepresenting workers in South Asia, Southeast Asia and Latin America. While translation support was extensive, some linguistic and cultural nuance may have been lost in non-English interviews. Despite these limitations, the multi-sited design, theoretically informed sampling and explicit reflexivity contribute methodological rigor to AI labour research, offering transferable insights for understanding invisibility across AI supply chains.

4. Findings: Dimensions of Invisibility in AI Labour

Building on the three-dimensional framework established in Section 2.4, this section presents empirical findings from 80 interviews with workers across data annotation, content moderation and RLHF contracting. The analysis reveals three interconnected dimensions through which invisibility operates: material structures that obscure labour, epistemic hierarchies that devalue worker knowledge and narrative framings that discursively erase human contributions. These findings ground EVF empirically, demonstrating how each dimension operates in practice across diverse regional contexts [53].

4.1. Material Invisibility: Structural Precarity and Economic Erasure

The most consistent theme was structural invisibility—workers rendered invisible not only to end-users but within corporate hierarchies that fragment and outsource labour. A Kenyan data annotator stated: *"We were told we are part of something innovative, but nobody sees us. The engineers get the credit, the managers get promotions, and we just disappear into the background"* (P17, Kenya). This invisibility was materially enforced through contractual precarity, wage opacity and geographic displacement.

Contractual arrangements systematically obscured workers' relationships to AI systems they built. Of 68 participants in platform roles, 61 worked through intermediary firms with no direct contract to end clients. A Nigerian content moderator explained: *"I moderate for a major platform, but technically I work for an outsourcing company in Lagos. If something goes wrong, the platform says, 'that's your employer's problem', and my employer says 'we just follow platform rules'. Nobody is responsible"* (P23, Nigeria). This double displacement enabled AI companies to externalise labour costs and liability while maintaining plausible deniability about working conditions.

Wage structures enforced invisibility through opacity and suppression. RLHF annotators in Kenya earned \$2-4 hourly while US-based contractors earned \$15-25 for identical work, yet these differentials were never explained. As a South African annotator noted: *"We don't know how they decide our rates. There's no transparency. You either accept it or you don't work"* (P34, South Africa). Workers described unpaid time waiting for tasks, completing training modules or handling technical glitches. *"They only pay for 'active' time, but I'm sitting here six hours waiting for tasks. That's my time too"* (P29, Ghana).

The psychological toll manifested as diminished self-worth. Workers internalised their erasure, questioning their value while recognising their centrality. *"Without us, there is no data. But the companies act as if they are doing us a favour by giving us these jobs"* (P52, Nigeria). This paradox—simultaneously essential and disposable—characterised workers' material condition across all sites.

4.2. Epistemic Invisibility: Knowledge Extraction and Expertise Denial

Beyond material marginalisation, workers faced systematic epistemic exclusion: their knowledge extracted while their status as knowledge producers was denied. This operated through expertise appropriation without recognition, testimonial silencing and exclusion from governance processes.

Many participants expressed frustration at being reduced to "button pushers" despite bringing substantial domain expertise. An RLHF annotator with a master's degree in linguistics described: *"I'm making judgments about language nuance,* *cultural appropriateness, contextual meaning. This is skilled cognitive work. But to them, we are just cheap labour clicking buttons"* (P44, South Africa). A Nigerian content moderator added: *"When we flag content, it is because we know how people really speak here—the slang, the context, the cultural references. But the guidelines come from somewhere else, from Silicon Valley. They don't listen to us"* (P23, Nigeria).

This expertise denial had consequences for AI system quality. Several workers described providing feedback about cultural misinterpretations or biased guidelines, only to have input ignored. "We know these models better than the engineers sometimes—we see every output, every error pattern. But no one asks us how they should be improved" (P38, South Africa). When workers raised concerns, they faced retaliation risks. "If you complain, they can cut your contract. So, you learn to be quiet" (P58, UAE). This created what Fricker [17, 21] terms "testimonial injustice": workers' knowledge systematically discredited based on social position rather than content validity.

The ultimate form of epistemic invisibility was workers' exclusion from "AI ethics" discourse. Multiple participants noted the irony that ethics discussions focused on algorithmic bias while ignoring labour conditions. "*I see all these conferences about responsible AI, fairness, transparency. But they never talk about us—the people doing the work"* (P15, Kenya). A South African annotator made the connection explicit: *"They debate whether AI is biased against minorities. Meanwhile, we are the minorities doing the annotation and nobody asks about our working conditions. That's the real bias"* (P34, South Africa).

4.3. Narrative Invisibility: Discursive Erasure and Automation Mythology

The third dimension concerned how workers were written out of dominant narratives about AI through automation discourse, media representations celebrating "smart machines" while omitting human trainers and workers' own internalisation of invisibility.

Several participants described encountering headlines about AI achievements and recognising their labour behind them yet seeing no acknowledgment. *"I see headlines about AI beating doctors or lawyers. And I think, 'Wait—that's my data behind it, my annotations, my judgments'. But no one mentions us. It's like we never existed"* (P34, South Africa). Another worker reflected: *"The headlines say, 'AI detected hate speech', but it was me sitting here twelve hours a day reviewing the worst content imaginable to teach the model. Where am I in that story?"* (P17, Kenya).

This discursive erasure extended beyond media to corporate communications. Platform transparency reports detailed "AI-powered content moderation" without mentioning human moderators. Model cards described "human feedback" without specifying labour conditions under which feedback was extracted. "*They write about 'human-in-the-loop' like we're just a technical component, not actual people with lives and families"* (P52, Nigeria).

Workers described how automation mythology shaped public perception of their work. Several reported family members not understanding their jobs or viewing them as temporary *"until AI takes over". A Kenyan annotator explained: "My parents don't understand what I do. They ask when I'll get a 'real job'. They think what I do is temporary, that AI will replace me. But I'm the one teaching the AI"* (P17, Kenya).

Some workers explicitly connected narrative invisibility to colonial patterns. A South African participant reflected: *"This is just like the colonial economy—we provide the raw material, the knowledge, the labour and they take it North and put their brand on it. Then they tell the world they invented something. We disappear from the story completely"* (P38, South Africa).

4.4. Synthesis: The Ethical Visibility Framework in Practice

These findings demonstrate that invisibility operates as a system: material structures enable epistemic extraction, which is naturalised through narrative erasure, which in turn justifies continued material exploitation. Workers are not merely absent from view but actively obscured through interlocking mechanisms that benefit AI companies economically, epistemically and politically[3].

Yet visibility is not absolute. Workers resist through informal collectives and work slowdowns. *"We made our own WhatsApp group. If someone doesn't get paid, we warn the rest. If there's a trick to do tasks faster, we share it. They try to keep us isolated, but we find each other"* (P17, Kenya). These micro-resistances reveal that invisibility is contested terrain, not total domination.

Table 2. *Empirical Manifestations of the Ethical Visibility Framework.*

EVF Dimension	Structural Mechanism	Empirical Finding	Worker Quote
Material Invisibility	Platform intermediation	61 of 68 platform workers employed through outsourcing firms with no direct client relationship.	"I moderate for a major platform, but technically I work for an outsourcing company... Nobody is responsible" (P23, Nigeria).
	Wage suppression	\$2-4/hour (Kenya) vs. \$15-25/hour (US) for	"They only pay for 'active' time, but I'm sitting

EVF Dimension	Structural Mechanism	Empirical Finding	Worker Quote
		identical RLHF work; undocumented unpaid waiting time.	here six hours waiting for tasks" (P29, Ghana).
	Geographic displacement	81% of sample (65 of 80) based in Global South performing work for Northern companies.	"Without us, there is no data. But companies act like they're doing us a favour" (P52, Nigeria).
Epistemic Invisibility	Expertise appropriation	Workers with domain knowledge (linguistics, culture, context) reduced to "button pushers".	"I'm making judgments about language nuance... But to them, we are just cheap labour clicking buttons" (P44, South Africa).
	Testimonial silencing	Feedback channels exist but carry retaliation risks; input systematically ignored	"If you complain, they can cut your contract. So, you learn to be quiet" (P58, UAE).
	Governance exclusion	Zero participants invited to participate in company AI ethics processes despite centrality to outcomes.	"All these conferences about responsible AI... But they never talk about us" (P15, Kenya).
Narrative Invisibility	Automation mythology	Media and corporate communications credit "AI" for outputs produced through extensive human labour.	"Headlines say 'AI detected hate speech', but it was me here twelve hours reviewing content" (P17, Kenya).
	Transitional framing	Workers' roles described as temporary, low-skill, awaiting automation—despite no evidence of replacement.	"My parents think what I do is temporary, that AI will replace me. But I'm teaching the AI" (P17, Kenya).
	Colonial patterns	Global South labour extracted while innovation credited to Global North companies.	"They take our labour North and put their brand on it. We disappear from the story" (P38, South Africa).

5. Discussion: The Ethical Visibility Framework and Its Implications

The findings demonstrate that invisibility in AI labour operates as an integrated system rather than isolated incidents. As the empirical synthesis in Section 4.4 demonstrates, workers are structurally marginalised through outsourcing, epistemically excluded from knowledge production and narratively erased from dominant AI imaginaries. This discussion elaborates the Ethical Visibility Framework's theoretical contributions and outlines policy implications for making labour visible.

5.1. The Ethical Visibility Framework: Conceptual Architecture and Theoretical Integration

The findings demonstrate that invisibility in AI labour operates as an integrated system rather than isolated incidents. As the empirical synthesis in Section 4.4 demonstrates, workers are structurally marginalised through outsourcing, epistemically excluded from knowledge production and narratively erased from dominant AI imaginaries[56]. This discussion elaborates the Ethical Visibility Framework's theoretical contributions and outlines policy implications for making labour visible.

Material invisibility operates through platform intermediation, supply chain fragmentation and contractual displacement. As Miceli and Posada [34] argue, the "data-production dispositive" systematically organises data workers while rendering their contributions invisible. Our findings extend this by documenting how 90% of platform workers (61 of 68) had no direct contractual relationship with the AI systems they built, enabling companies to externalise liability. Geographic displacement compounds this: 81% of participants laboured in the Global South for Northern companies, reproducing colonial patterns (Couldry & Mejias [13, 18]). Wage suppression—$2-4 hourly in Kenya versus $15-25 for identical RLHF work in the US—was enabled through this structural invisibility.

Epistemic invisibility functions through expertise hierarchies and knowledge extraction. Building on Sambasivan and Veeraraghavan [46] analysis of "deskilling", EVF shows how workers' domain knowledge is systematically harvested while their status as knowledge producers is denied. This operates as what Fricker [21, 35] terms "epistemic injustice": workers' testimonies are systematically discredited based on social position rather than content validity. Our findings document zero instances of workers being invited to participate in company AI ethics processes despite their centrality to system outcomes.

This exclusion is strategic: recognising workers as epistemic contributors would undermine wage suppression and demand participation in governance [58].

Narrative invisibility works through automation mythology and governance frameworks that attribute agency to "AI" while erasing human trainers. As Suchman [52] argues, imaginaries of autonomous AI are performative—they make certain futures thinkable by obscuring present human work. Our findings show workers encountering media headlines about "AI achievements" and recognising their labour yet seeing no acknowledgment. This extends to regulatory frameworks: the EU AI Act foregrounds algorithmic risk while neglecting labour conditions, reproducing invisibility at the governance level [24].

As Section 4.4 demonstrates empirically, these dimensions constitute a self-sustaining cycle that requires coordinated intervention across all three levels simultaneously.

[KEEP DIAGRAM EXACTLY AS IS - INCLUDING INTERVENTION POINTS]

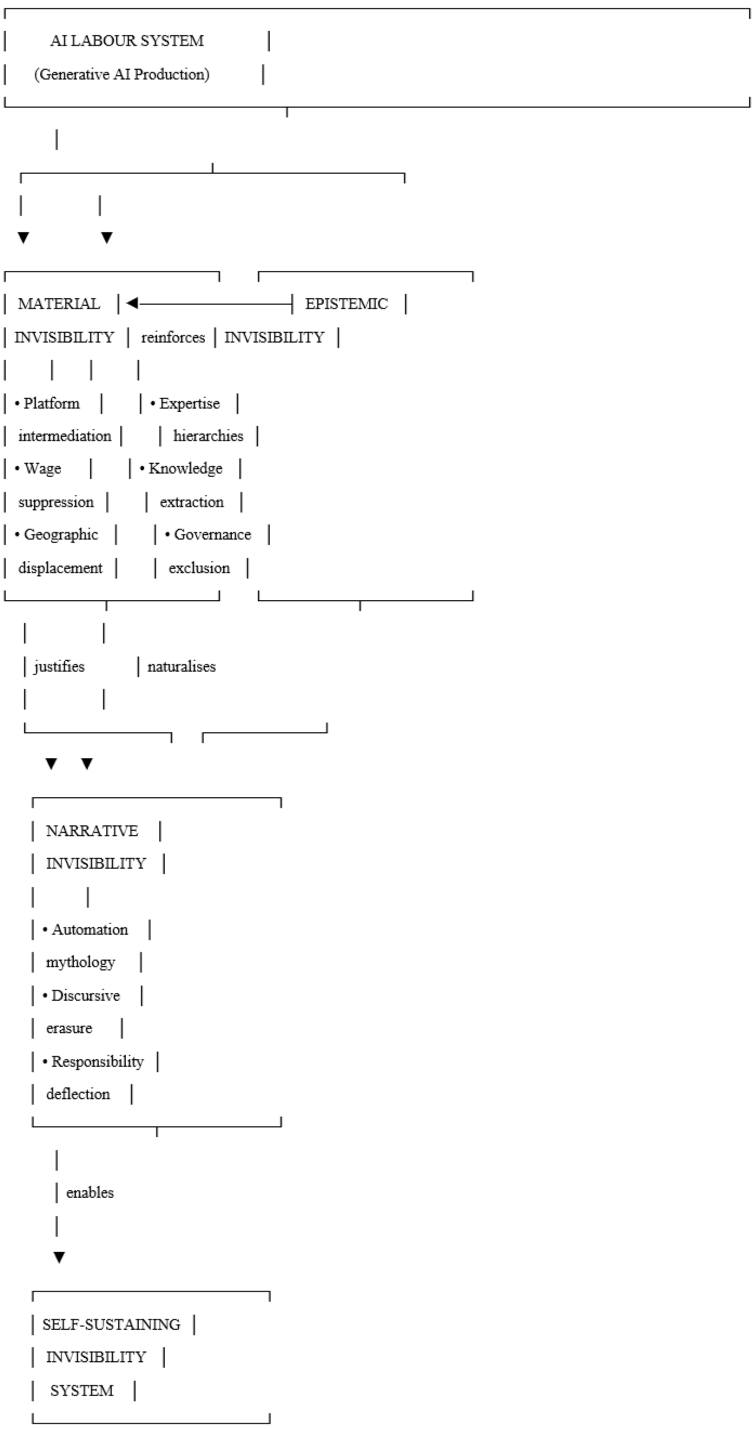

Figure 1. *The Ethical Visibility Framework—Reinforcing Dimensions of Invisibility.*

INTERVENTION POINTS (breaking the cycle):
MATERIAL LEVEL:
1) Supply chain disclosure requirements
2) Labour protections in AI regulation
3) Direct contracting models
4) Wage transparency mandates
EPISTEMIC LEVEL:
1) Worker participation in AI governance
2) Co-authorship credit for datasets
3) Expertise recognition frameworks
4) Inclusion in "responsible AI" processes
NARRATIVE LEVEL:
1) Labour documentation in model cards
2) Worker acknowledgment in media coverage
3) Counter-narratives challenging automation mythology
4) Corporate transparency reporting on human contributions

EVF's theoretical contribution lies in demonstrating how these dimensions interact rather than treating labour conditions, knowledge politics and discourse as separate domains. This integrative approach extends Gray and Suri [23] "ghost work" concept by theorising invisibility's structural mechanisms; advances Miceli et al.'s (2022) call to study power rather than bias by centring labour relations; and operationalises postcolonial critiques [4] through empirical documentation of epistemic extraction.

5.2. Contributions to Critical AI Ethics and Labour Scholarship

EVF makes three primary contributions: it extends recent work mapping AI labour (Catanzariti et al. [10]; Lampinen et al. [31]) and documenting precarity (Gray & Suri [23]; Roberts [43]) by providing systematic framework for analysing how invisibility operates across multiple dimensions simultaneously. Where existing work documents labour conditions descriptively, EVF demonstrates how material, epistemic and narrative invisibility constitute an integrated system requiring coordinated intervention.

Second, EVF advances postcolonial and decolonial critiques of AI (Birhane [4]; Mohamed et al. [38]) by empirically grounding theoretical arguments about epistemic extraction. The framework shows how Global South workers situated knowledge about language, culture and context is systematically harvested while credit accrues to Northern institutions—a pattern Couldry and Mejias [13] term "data colonialism" but which EVF specifies operates through invisibility mechanisms at structural, epistemic and discursive levels.

Third, EVF extends critical political economy approaches (Srnicek [49]; Zuboff [62] by theorising cognitive and affective extraction alongside economic exploitation. The framework positions invisibility not as an incidental feature of platform capitalism but as a necessary condition for AI's current mode of production. If annotators were recognised as epistemic contributors, their labour would command higher wages and participation rights; if content moderators' trauma were acknowledged, companies would face pressure for psychological support. Invisibility thus enables continued extraction by preventing such recognition and its attendant obligations.

5.3. Policy and Governance Implications: Making Labour Visible

Making AI labour visible requires interventions at multiple scales, targeting each dimension while recognising their interconnections. These recommendations emerge from worker testimonies and analytical findings.

Material interventions must address supply chain opacity and contractual fragmentation. AI companies should implement supply chain disclosure requirements analogous to conflict minerals reporting, documenting labour conditions throughout annotation and moderation supply chains. Concretely, this should take the form of a mandatory annual AI Labour Supply Chain Report — modelled on obligations under the UK Modern Slavery Act 2015 — requiring companies above a defined revenue threshold to disclose all third-party annotation and moderation contractors, worker volumes, pay ranges, and jurisdictions of employment. The EU AI Act could be amended to include labour impact assessments as part of high-risk AI system evaluation, treating worker precarity as a form of systemic risk alongside algorithmic bias. Non-compliance with such assessments should carry the same market access consequences as algorithmic non-conformity, ensuring enforcement teeth that voluntary corporate commitments have consistently failed to provide, and would align with long-standing international labour standards [27].

Epistemic interventions demand recognising data workers as knowledge producers with legitimate expertise. This could take several forms: co-authorship credit on papers using datasets workers created; mandatory consultation with annotation teams during model development; inclusion of worker representatives on corporate AI ethics boards with decision-making power; and participatory design processes that position workers as co-researchers. To operationalise these principles, academic journals and dataset repositories should adopt a mandatory Labour Contribution Statement as a condition of submission — analogous to existing data availability statements — requiring authors to specify worker compensation, geographic distribution, and organisational relationships. Professional bodies such as the ACM and IEEE are well positioned to develop and enforce such standards through publication eligibility requirements. As DiSalvo et al. [19] argue, workplace democracy in data work requires first making work socially visible, then creating mechanisms for workers to exercise collective power.

Narrative interventions require linguistic and documentary shifts across multiple sites. Research papers should specify worker conditions in methods sections; corporate communications should acknowledge human training labour; news coverage should interview annotators alongside developers; and model cards should include labour documentation sections. The existing model card standard should be extended to include a mandatory Human Labour section — specifying worker numbers, geographic base, compensation structure, and working conditions — with major repositories such as Hugging Face and Zenodo rejecting submissions that omit it. Science and technology journalists, meanwhile, should treat worker testimony as a standard sourcing requirement in AI capability reporting, a norm that professional bodies could formalise through updated editorial guidance. These changes would challenge automation mythology by making human contributions legible, enabling public accountability for labour conditions.

As Section 4.4 demonstrates empirically, these dimensions constitute a self-sustaining cycle that requires coordinated intervention across all three levels simultaneously.

Worker testimonies underscore implementation urgency. As one participant asked: "Why can't they just say that people are behind the AI? Why hide us?" (P44, South Africa). This question cuts to the core ethical failure: AI systems are presented as autonomous precisely to obscure human costs. Governance frameworks that treat AI ethics as technical optimisation while ignoring labour conditions will inevitably reproduce the invisibilities they claim to address. Making labour visible is not an auxiliary concern but a foundational ethical requirement for legitimate AI development.

6. Conclusion: Toward Ethically Visible AI Labour

This study has argued that the invisibility of human labour in generative AI is not an oversight but a structural feature of contemporary AI production. Through the Ethical Visibility Framework, we have demonstrated how material, epistemic and narrative mechanisms work together to obscure the workers who train, align and moderate AI systems. Drawing on 80 interviews across eight countries, the analysis reveals that this invisibility is both globally patterned—concentrating precarity in the Global South—and locally experienced through specific practices of wage suppression, expertise denial and discursive erasure.

EVF contributes to AI ethics scholarship by providing a systematic framework for analysing how invisibility is produced and maintained across multiple dimensions simultaneously. It extends digital labour studies by foregrounding epistemic and narrative dimensions alongside material conditions and advances postcolonial Science and Technology Studies by demonstrating how colonial extraction patterns persist in AI supply chains. Methodologically, the study challenges Northern epistemological dominance in AI ethics discourse by centring Global South workers' experiences and analytical insights as legitimate knowledge rather than supplementary testimony.

For practitioners and policymakers, EVF identifies concrete intervention points. Material visibility requires supply chain disclosure requirements and labour protections integrated into AI regulation; epistemic visibility demands worker participation in AI governance through co-authorship credit, consultation mechanisms and ethics board representation; and narrative visibility necessitates documentary practices that acknowledge human labour in research papers, corporate communications and media coverage. Technology companies, regulators and researchers each have roles in making AI labour visible—not as a corporate social responsibility gesture but as an ethical prerequisite for legitimate AI development.

Future research should extend EVF across different AI sectors, examine worker organising strategies in greater depth and develop participatory methodologies that position workers as co-researchers. Comparative work across additional Global South regions would illuminate regional variations in labour conditions and invisibility regimes.

Ultimately, ethical AI cannot be built on invisible labour. The workers who teach ChatGPT to converse, who moderate harmful content to train safety systems and who label images to enable computer vision are not peripheral service providers but epistemic agents whose contributions make AI possible [30]. Recognising their visibility is not merely a matter of justice—it is a precondition for AI systems that reflect human values rather than obscure human costs. Workers' invisibility is strategic, not natural; their marginalisation is engineered, not inevitable. Challenging this requires treating visibility not as an add-on but as a foundational ethical commitment—recognising that those who build AI systems deserve recognition, protection and participation in shaping their futures.

Abbreviations

AI	Artificial Intelligence
CSAM	Child Sexual Abuse Material
ESRC	Economic and Social Research Council
EU	European Union
EVF	Ethical Visibility Framework
LLM	Large Language Model
NDA	Non-Disclosure Agreement
OECD	Organisation for Economic Co-operation and Development
RLHF	Reinforcement Learning from Human Feedback
STS	Science and Technology Studies

Author Contributions

Achi Iseko: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing

Conflicts of Interest

The author declares no conflicts of interest.

References

[1] Abdalla, M. (2021). The Grey Hoodie Project: Big Tobacco, Big Tech, and the threat on academic integrity. *Big Data & Society*, 8(1), 1-13. https://doi.org/10.1177/20539517211020675

[2] Abebe, R., Barocas, S., Kleinberg, J., Levy, K., Raghavan, M., & Robinson, D. (2020). Roles for computing in social change. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAccT '20)*, 252-260. Association for Computing Machinery. https://doi.org/10.1145/3351095.3372871

[3] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv: 1606.06565.

[4] Birhane, A. (2020). Algorithmic colonization of Africa. *Scripted*, 17(2), 389-409. https://doi.org/10.2966/scrip.170220.389

[5] Birhane, A. (2021). Algorithmic injustice: A relational ethics approach. *Patterns*, 2(2), 100205. https://doi.org/10.1016/j.patter.2021.100205

[6] Birhane, A., & Guest, O. (2023). Towards decolonising computational sciences. *Women, Gender & Research*, 31(2), 60-73. https://doi.org/10.7146/kkf.v31i2.133153

[7] Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.

[8] Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101. https://doi.org/10.1191/1478088706qp063oa

[9] Casilli, A., & Posada, J. (2019). The platformization of labor and society. In M. Graham & W. H. Dutton (Eds.), *Society and the internet: How networks of information and communication are changing our lives* (pp. 293-306). Oxford University Press.

[10] Catanzariti, B., Chandhiramowuli, S., Mohamed, S., Natarajan, S., Prabhat, S., Raval, N., Taylor, A. S., & Wang, D. (2021). The Global Labours of AI and Data Intensive Systems. In *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing* (CSCW '21 Companion), 319-322. https://doi.org/10.1145/3462204.3481725

[11] Charmaz, K. (2014). *Constructing grounded theory* (2nd ed.). SAGE Publications.

[12] Collins, P. H. (2000). *Black feminist thought: Knowledge, consciousness, and the politics of empowerment* (2nd ed.). Routledge.

[13] Couldry, N., & Mejias, U. A. (2019). *The costs of connection: How data is colonizing human life and appropriating it for capitalism*. Stanford University Press.

[14] Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

[15] Crawford, K., Dobbe, R., Kaziunas, E., Mathur, V., Richardson, R., Reisman, D., Schultz, J., Myers West, S., & Whittaker, M. (2019). AI and ethics in practice: The case for worker-centred design. AI Now Institute. https://ainowinstitute.org/publication/ai-and-ethics-in-practice

[16] Crenshaw, K. (1991). Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford Law Review*, 43(6), 1241-1299. https://doi.org/10.2307/1229039

[17] D'Ignazio, C., & Klein, L. F. (2020). *Data feminism*. MIT Press.

[18] Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56-62. https://doi.org/10.1145/2844110

[19] DiSalvo, C., Rothschild, A., Schenck, L. L., Shapiro, B. R., & DiSalvo, B. (2023). When Workers Want to Say No: A View into Critical Consciousness and Workplace Democracy in Data Work. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1), 1-23. https://doi.org/10.1145/3579515

[20] Fairclough, N. (2013). *Critical discourse analysis: The critical study of language* (2nd ed.). Routledge.

[21] Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.

[22] Graham, M., Hjorth, I., & Lehdonvirta, V. (2017). Digital labour and development: Impacts of global digital labour platforms and the gig economy on worker livelihoods. *Transfer: European Review of Labour and Research*, 23(2), 135-162. https://doi.org/10.1177/1024258916687250

[23] Gray, M. L., & Suri, S. (2019). *Ghost work: How to stop Silicon Valley from building a new global underclass*. Houghton Mifflin Harcourt.

[24] Greene, D., Hoffmann, A. L., & Stark, L. (2019). Better, nicer, clearer, fairer: A critical assessment of the movement for ethical artificial intelligence and machine learning. *Proceedings of the 52nd Hawaii International Conference on System Sciences*, 2122-2131. https://doi.org/10.24251/HICSS.2019.258

[25] Haraway, D. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3), 575-599. https://doi.org/10.2307/3178066

[26] Harding, S. (1991). *Whose science? Whose knowledge? Thinking from women's lives*. Cornell University Press.

[27] ILO (International Labour Organization). (2008). *ILO Declaration on Social Justice for a Fair Globalization*. International Labour Office.

[28] Irani, L. (2015). Difference and dependence among digital workers: The case of Amazon Mechanical Turk. *South Atlantic Quarterly*, 114(1), 225-234. https://doi.org/10.1215/00382876-2831665

[29] Jasanoff, S. (2015). Future imperfect: Science, technology, and the imaginations of modernity. In S. Jasanoff & S.-H. Kim (Eds.), *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power* (pp. 1-33). University of Chicago Press.

[30] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097-1105.

[31] Lampinen, A., Comber, R., Chandhiramowuli, S., Møller, N. H., & Taylor, A. S. (2024). The Work of AI: Mapping Human Labour in the AI Pipeline. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing* (CSCW Companion '24), 728-731. https://doi.org/10.1145/3678884.3681825

[32] Markham, A. N. (2018). Afterword: Ethics as methods, methods as ethics. In U. Flick (Ed.), *The SAGE handbook of qualitative data collection* (pp. 537-553). SAGE Publications.

[33] Mbembe, A. (2001). *On the postcolony*. University of California Press.

[34] Miceli, M., & Posada, J. (2022). The Data-Production Dispositif. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2), 460: 1-460: 37. https://doi.org/10.1145/3555561

[35] Miceli, M., Posada, J., & Yang, T. (2022). Studying Up Machine Learning Data: Why Talk About Bias When We Mean Power? *Proceedings of the ACM on Human-Computer Interaction*, 6(GROUP), 1-14. https://doi.org/10.1145/3492853

[36] Mhlambi, S. (2020). From rationality to relationality: Ubuntu as an ethical and human rights framework for artificial intelligence governance. *Carr Centre for Human Rights Policy*, Harvard Kennedy School. https://carrcenter.hks.harvard.edu/files/cchr/files/ccdp_2020-009_sabelo_mhlambi.pdf

[37] Milan, S., & Treré, E. (2019). Big data from the South(s): Beyond data universalism. *Television & New Media*, 20(4), 319-335. https://doi.org/10.1177/1527476419837739

[38] Mohamed, S., Png, M.-T., & Isaac, W. (2020). Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, 33(4), 659-684. https://doi.org/10.1007/s13347-020-00405-8

[39] Ndlovu-Gatsheni, S. J. (2015). Decoloniality as the future of Africa. *History Compass*, 13(10), 485-496. https://doi.org/10.1111/hic3.12264

[40] Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.

[41] Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools, and societies*. Princeton University Press.

[42] Pasquinelli, M., & Joler, V. (2021). The nooscope manifested: AI as instrument of knowledge extractivism. *AI & Society*, 36(4), 1263-1280. https://doi.org/10.1007/s00146-021-01280-2

[43] Roberts, S. T. (2019). *Behind the screen: Content moderation in the shadows of social media*. Yale University Press.

[44] Russell, S., & Norvig, P. (2020). *Artificial intelligence: A modern approach* (4th ed.). Pearson.

[45] Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. M. (2021). "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1-15. https://doi.org/10.1145/3411764.3445518

[46] Sambasivan, N., & Veeraraghavan, R. (2022). The Deskilling of Domain Expertise in AI Development. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1-14. https://doi.org/10.1145/3491102.3517578

[47] Seaver, N. (2017). Algorithms as culture: Some tactics for the ethnography of algorithmic systems. *Big Data & Society*, 4(2), 1-12. https://doi.org/10.1177/2053951717738104

[48] Simonite, T. (2021, April 6). Millions of people train AI systems—without knowing it. *Wired*. https://www.wired.com/story/millions-people-train-ai-systems-without-knowing-it/

[49] Srnicek, N. (2017). *Platform capitalism*. Polity Press.

[50] Star, S. L. (1991). Power, technology and the phenomenology of conventions: On being allergic to onions. In J. Law (Ed.), *A sociology of monsters: Essays on power, technology and domination* (pp. 26-56). Routledge.

[51] Star, S. L., & Strauss, A. (1999). Layers of silence, arenas of voice: The ecology of visible and invisible work. *Computer Supported Cooperative Work*, 8(1-2), 9-30. https://doi.org/10.1023/A:1008651105359

[52] Suchman, L. (2007). *Human-machine reconfigurations: Plans and situated actions* (2nd ed.). Cambridge University Press.

[53] Suchman, L. (2020). Algorithmic warfare and the reinvention of accuracy. *Critical Studies on Security*, 8(2), 175-187. https://doi.org/10.1080/21624887.2020.1760587

[54] Tubaro, P., Casilli, A. A., & Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1), 1-12. https://doi.org/10.1177/2053951720919776

[55] Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433-460. https://doi.org/10.1093/mind/LIX.236.433

[56] van Dijk, T. A. (1993). Principles of critical discourse analysis. *Discourse & Society*, 4(2), 249-283. https://doi.org/10.1177/0957926593004002006

[57] Wallace, R., & Mhlambi, S. (2022). Centering African values in artificial intelligence research: A call for collaboration. *AI and Ethics*, 2(4), 689-696. https://doi.org/10.1007/s43681-022-00156-7

[58] West, S. M., Whittaker, M., & Crawford, K. (2019). *Discriminating systems: Gender, race, and power in AI*. AI Now Institute. https://ainowinstitute.org/discriminatingsystems.html

[59] Whittaker, M., Crawford, K., Dobbe, R., Fried, G., Kaziunas, E., Mathur, V., West, S. M., Richardson, R., Schultz, J., & Schwartz, O. (2018). *AI Now Report 2018*. AI Now Institute. https://ainowinstitute.org/AI_Now_2018_Report.pdf

[60] Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56-75. https://doi.org/10.1177/0950017018785616

[61] Yudkowsky, E. (2008). Artificial intelligence as a positive and negative factor in global risk. In N. Bostrom & M. M. Ćirković (Eds.), *Global catastrophic risks* (pp. 308-345). Oxford University Press.

[62] Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.