"Artificial But Not Intelligent": the Essence of Ground-Level Artificial Intelligence (AI) Labor Experience and the Ontology of Intelligence

Brian Bantugan
St. Paul University Manila

Abstract: *This study investigated the critical yet often invisible role of human labor in the development of artificial intelligence (AI). Utilizing an interpretive critical phenomenological approach, the research explored the lived experiences of ground-level workers—often termed "ghost workers"—who performed essential tasks such as data labeling, content moderation, and annotation. By integrating textual analysis of an investigative documentary set in Nairobi, Kenya, with a comprehensive literature review, the paper revealed how AI's celebrated "intelligence" was actually a derivative, technically mediated projection of human judgment.The findings highlighted significant ethical and systemic issues, including extreme economic precarity, psychological trauma from exposure to harmful content, and deceptive platform practices that limited worker autonomy. The research argued that these workers did not merely "input data" but constructed situated, tacit, and moral knowledge that was indispensable to AI functionality. However, this human element was systematically erased through aggregation, resulting in "polished" outputs that lacked the ethical reflexivity and contextual awareness of their human creators.Ultimately, the study challenged the metaphor of AI as a mirror of human intelligence, proposing instead that it functioned as a "refractive apparatus" that bent human cognition toward corporate priorities while obscuring labor exploitation. It concluded that the development of ethical and high-quality AI remained impossible without a radical transformation of the labor conditions and knowledge pipelines that sustained it.*

Keywords: *Ground-level AI workers, AI, Ethics, Human Intelligence, Epistemic Injustice*

I. Introduction

Artificial intelligence (AI) is widely celebrated as a transformative technology capable of automating complex tasks, enhancing decision-making, and revolutionizing industries. Media narratives often depict AI as autonomous, intelligent, and self-sufficient, suggesting a future in which human labor is increasingly redundant (Floridi, 2019; Russell & Norvig, 2021). Yet, beneath the polished interfaces and confident outputs of AI systems lies a hidden workforce whose labor makes these technologies possible. Individuals engaged in content labeling, moderation, and data annotation—often termed "ghost workers"—perform essential tasks under precarious conditions, producing the interpretive, contextual, and moral judgments that AI systems rely upon (Gray & Suri, 2019; Rahman & Sultana, 2025).

Despite their critical role, these workers remain largely invisible in public discourse. Their labor is undervalued, underpaid, and often associated with microtasks that obscure cognitive and emotional effort. As Gray and Suri (2019) note, AI's apparent intelligence is derivative, reflecting human judgment encoded at scale rather than autonomous reasoning. Workers report enduring psychological harm, moral injury, and coercive labor conditions, highlighting the ethical and human costs of AI development (Gonzalez & Matias, 2025; Time, 2023). These experiences underscore the disconnect between the technical abstraction of AI and the embodied, affective labor that sustains it.

Crucially, the knowledge produced by these workers is epistemically significant. It is situated, tacit, and moral, encompassing experiential insights into human behavior, social norms, and ethical dilemmas embedded in AI outputs (Polanyi, 1966; Haraway, 1988). Workers' collective resistance and advocacy, manifested through networks and associations, further transform private suffering into public, transformative knowledge that challenges dominant narratives of innovation and neutrality (MDPI, 2025). In this light, AI is not simply a technological artifact but a socio-technical system deeply shaped by human intelligence, labor exploitation, and structural inequality.

This study investigated the experiences of ground-level AI workers and the knowledge they generate, addressing critical questions about visibility, economic precarity, psychological harm, coercion, and moral agency. By centering human labor in discussions of AI development, the research illuminates the ethical, epistemic, and social dimensions of AI production, revealing that the prospects for ethical and high-quality AI are inextricably tied to the conditions under which knowledge is produced (Crawford, 2021; Gray & Suri, 2019). In doing so, the study challenges conventional metaphors of AI as a mirror of human intelligence, arguing instead that AI functions as a refractive apparatus that reshapes and redistributes human cognition according to structural priorities (Crawford, 2021).

Ultimately, understanding AI requires more than technical analysis; it demands attention to the human experiences, labor dynamics, and moral complexities that underlie the technologies we increasingly rely upon. By foregrounding these issues, this research contributes to a more nuanced, ethical, and human-centered discourse on AI and its future.

Invisibility and "Ghost Work"

A growing body of research highlights how human labor underlies the functioning of AI systems while remaining hidden in public narratives about autonomous technologies. Gray and Suri's foundational work on *ghost work* describes the invisible, task-based labor—such as labeling, data cleaning, and content moderation—that sustains AI development yet receives little recognition or credit (Gray & Suri, 2019). This hidden workforce is often outsourced through digital platforms and framed as automated or microtask work, obscuring the human agency involved in meaning-making and error correction (Gray & Suri, 2019; Rahman & Sultana, 2025). Scholars note that such labor is undervalued and hidden by design, sustaining the illusion that AI systems operate independently of human input (Just Tech, 2025; Time, 2023).

Economic Precarity and Platform Labor

Digital platform labor used in AI development is structurally precarious. Workers frequently earn below living wages, lack employment protections, and are classified as independent contractors, which erodes access to benefits and legal recourse (Time, 2023; Simet, 2025). Platform algorithmic management intensifies instability by controlling task allocation, performance metrics, and continued access to work through opaque systems, leaving workers dependent on platforms for income while holding minimal autonomy (SumiNote, 2025; Casilli, 2024). These dynamics reproduce global labor inequalities, as low-wage workers in the Global South disproportionately supply labor critical to AI training. Such conditions reflect broader patterns of data colonialism and labor exploitation embedded in digital economies (Casilli, 2024; Simet, 2025).

Psychological Harm and Affective Labor

Exposure to harmful content and the emotional labor of digital work have documented psychological consequences. Systematic reviews of content moderation work find that prolonged exposure to violent, sexual, or abusive material correlates with lasting emotional distress, with existing mental health measurement tools insufficiently adapted to this context (Gonzalez & Matias, 2025). Related literature on algorithmic control in gig work highlights how monitoring and performance pressure can contribute to burnout, stress, and diminished emotional well-being, suggesting that psychological strain is a significant hazard of platform labor (BMC Psychology, 2023; Safety Science, 2023). These insights contextualize how emotional and cognitive burdens are inherent to tasks that shape AI outputs.

Worker Agency, Resistance, and Collective Identity

Despite structural constraints, studies show that gig workers develop forms of agency and collective identity through networks and advocacy. Research on digital labor platforms identifies how workers build social networks to negotiate power imbalances and resist exploitative conditions, challenging opaque algorithmic governance and reinforcing their role in shaping labor conditions (MDPI, 2025). Such collective responses echo broader calls for recognition of digital workers' rights and for more equitable governance in platform economies.

Epistemic Contributions and AI Knowledge Production

Critical scholarship on AI ethics and labor underscores that human workers generate indispensable knowledge that undermines dominant depictions of AI as neutral or solely machine-driven. Human labeling and moderation are interpretive tasks that embed contextual judgments into data used for training machine learning models, making AI outcomes dependent on the situated, tacit knowledge of laborers (Gray & Suri, 2019; Just Tech, 2025). These insights reflect how epistemic contributions from marginalized workers are essential to AI performance yet marginalized in technical discourses, reinforcing the argument that ethical AI must engage with labor conditions and the quality of human-generated data.

The Hidden Workers Behind Artificial Intelligence | FULL DOCUMENTARY

Set against the backdrop of Nairobi, Kenya, this 40-minute investigative documentary by 7 News Spotlight (https://www.youtube.com/watch?v=z_01q3boQ6c) uploaded June 16, 2025 in YouTube (with 1.15 million subscribers) exposes the hidden human labor sustaining the global artificial intelligence (AI) revolution. While AI is widely celebrated for its speed, convenience, and transformative promise—from everyday digital tasks to future medical breakthroughs—the film reveals the largely invisible workforce that makes these systems possible: African data labelers who train AI by interpreting, categorizing, and moderating vast amounts of data. Through the personal stories of Joan, Michael, Chi-Chi, and other workers, the documentary uncovers a system marked by extreme precarity, opaque contracting practices, racialized pay disparities, and repeated exposure to psychologically damaging content, including medical imagery, extreme violence, and child exploitation.

The film traces how global tech companies outsource critical AI labor through intermediary firms operating within the rapidly expanding gig economy, enabling low-cost, flexible work while diffusing accountability. It documents allegations of non-payment, deceptive task design, intrusive data requests involving children and families, and minimal mental health support, raising urgent ethical and legal concerns that echo conditions associated with modern slavery. Corporate assurances of ethical sourcing and compliance are juxtaposed with workers' lived realities, external evaluations, and confrontations with company leadership that reveal gaps in transparency and enforcement.

Concluding with the emergence of worker organizing efforts in Kenya, the documentary reframes AI not as an autonomous technological achievement but as a deeply human system built on unequal global labor relations. It calls on policymakers, corporations, and AI users worldwide to confront the moral costs embedded in everyday AI use and to recognize ethical labor practices as foundational to the future of artificial intelligence.

II. Theoretical Framework

The exploration of ground-level AI workers, their constructed knowledge, and the ethical implications of artificial intelligence is grounded in **three interrelated theoretical perspectives**: **Feminist Standpoint Theory, Critical Political Economy, and Constructivist Learning Theory**. Together, these frameworks provide conceptual scaffolding to understand how labor, knowledge, and power intersect in the production of AI.

Feminist Standpoint Theory. Feminist standpoint theory posits that knowledge is socially situated and that marginalized or oppressed groups possess unique epistemic access to the structures of power that dominate society (Haraway, 1988; Harding, 2004). Applied to AI labor, this perspective highlights that ground-level workers—often underpaid, invisible, and located in economically disadvantaged regions—develop knowledge from their lived experiences that is epistemically valuable yet often unrecognized in dominant AI narratives. These workers construct **situated, tacit, and morally informed knowledge** through daily interaction with AI systems, offering insights into the human dimensions of technology that cannot be captured by purely technical or managerial perspectives. The emphasis on standpoint knowledge legitimizes workers' counter-narratives, showing that their experiences are essential for understanding the ethical, social, and cognitive realities of AI development.

Critical Political Economy. Critical political economy examines how economic structures, labor relations, and power asymmetries shape technological development and knowledge production (Fuchs, 2017; Mosco, 2009). Within the context of AI, this perspective illuminates how **global inequalities, precarious work conditions, and structural exploitation** influence both the labor of ground-level AI workers and the quality of data that feeds AI systems. AI platforms, designed to optimize efficiency and profitability, create conditions where labor is coerced, undervalued, and often psychologically harmful. The theory underscores that AI systems are not neutral technological artifacts; they are socio-technical constructs embedded with the biases, constraints, and inequalities of the labor that produces them. Critical political economy thus enables an understanding of AI as a **refractive apparatus**, in which human intelligence is shaped, filtered, and redistributed by economic and platform pressures.

Constructivist and Experiential Learning Theory. Constructivist learning theory asserts that knowledge is actively constructed through interaction with the environment, rather than passively absorbed (Piaget, 1970; Berger & Luckmann, 1966). Complementing this, experiential learning theory emphasizes that learning occurs through reflection on concrete experiences (Kolb, 1984). Applied to AI labor, these frameworks explain how workers develop **tacit, embodied, and context-sensitive knowledge** through repetitive tasks, sensemaking, and moral engagement with the content they annotate and moderate. Their learning is both cognitive and affective, integrating pattern recognition, ethical reflection, and anticipatory judgment to produce meaningful outputs that AI systems later aggregate. Constructivist perspectives also account for the emergence of **collective knowledge and transformative agency**, as workers reflect on shared experiences, organize collectively, and advocate for ethical AI practices.

Integration of Frameworks. By combining feminist standpoint theory, critical political economy, and constructivist learning theory, this study situates ground-level AI workers as both **knowledge producers and ethical stakeholders**. Feminist standpoint theory foregrounds the epistemic significance of marginalized labor;

critical political economy situates this labor within structural inequalities and platform dynamics; and constructivist theory explains how knowledge emerges from lived experience, reflection, and embodied practice. Together, these theories provide a lens to examine not only **what AI workers know** but also **how and why this knowledge is produced, constrained, and ethically consequential**. This integrated theoretical framework allows the study to interrogate the **visibility, moral agency, and transformative potential** of workers' experiences, connecting labor realities to broader questions of AI ethics, epistemology, and social justice.

III. Statement of the Problem

Artificial intelligence (AI) was frequently portrayed as autonomous and objective, yet its development relied heavily on ground-level human workers whose labor remained largely invisible. These workers performed essential tasks—labeling, annotating, and moderating data—under conditions of economic precarity, coercion, and minimal recognition. They were often exposed to psychologically harmful content, lacked informed consent, and faced structural exploitation, yet their lived experiences generated unique, tacit, and moral knowledge that challenged dominant narratives of AI.

The invisibility and marginalization of these workers not only affected their well-being but also shaped the quality, ethics, and epistemic validity of AI outputs. Human judgment and emotional labor were embedded in AI systems while remaining unacknowledged, raising questions about the prospects of developing high-quality and ethical artificial general intelligence (AGI) without structural transformation of the labor–knowledge pipeline.

This study aimed to explore the lived experiences of ground-level AI workers, the knowledge they constructed, and the ethical and systemic implications of their labor in the ontology intelligence in AI..

IV. Methodology

This study employed an interpretive critical phenomenological approach to explore the lived experiences of ground-level AI workers and the knowledge constructed through their labor. Phenomenology focuses on the essence of lived experience, seeking to understand how individuals perceive and make meaning of phenomena in their daily lives (van Manen, 2014). The critical dimension of this approach emphasizes the interrogation of power, inequality, and structural constraints that shape these experiences (Finlay, 2011). By combining interpretive and critical lenses, this study sought to reveal not only the subjective realities of AI workers but also the broader socio-economic and technological conditions that mediate their labor, knowledge production, and ethical concerns.

Data Sources and Selection

Data were constructed from two primary textual sources: a purposively selected documentary video and a review of relevant literature. The documentary video ("The Hidden Workers Behind Artificial Intelligence | FULL DOCUMENTARY")**,** selected through purposive sampling, provided rich, first-hand accounts of AI workers' experiences, labor conditions, and moral dilemmas. Purposive selection ensured that the text chosen was directly relevant to the study's objectives and contained narratives representative of diverse geographic, socio-economic, and cultural contexts (Palinkas et al., 2015).

Complementing the video, the literature review served as a secondary textual source, allowing the researcher to contextualize the documentary data within existing scholarship on AI labor, epistemic extraction, ethical AI, and knowledge production. This dual-source strategy enabled triangulation, enhancing the trustworthiness and rigor of the study by cross-referencing lived accounts with theoretical and empirical insights (Creswell & Poth, 2018).

Data Construction Techniques

Textual analysis was the primary method for constructing data from the selected sources. The analysis involved a close reading of the documentary transcript, identifying salient passages, quotations, and narratives that reflected workers' experiences, ethical reflections, and tacit knowledge. Similarly, literature texts were analyzed for concepts, frameworks, and critical discussions relevant to AI labor, epistemic injustice, and moral injury.

Following textual analysis, textual synthesis was employed to integrate findings from both sources. This process involved identifying recurring themes, contrasts, and insights across sources, allowing for the construction of comprehensive interpretive patterns that connect individual experiences to broader structural and ethical issues. Through synthesis, the study generated new understanding of AI workers' knowledge as situated, embodied, and morally-informed, linking lived experience to critical theory and knowledge production debates (Braun & Clarke, 2013).

Ethical Considerations

Several ethical considerations were integral to the study design. First, the study respected the anonymity and confidentiality of workers represented in the documentary, ensuring that direct quotations did not compromise personal identity (Orb et al., 2001). Second, ethical reflexivity guided the interpretation of sensitive content, including exposure to trauma, exploitation, and coercion. The researcher acknowledged the affective weight of the data and the potential for secondary trauma in engagement with such material, adopting a careful, context-sensitive approach in analysis and reporting. Third, the study addressed epistemic ethics by recognizing the documentary subjects as legitimate knowledge producers, avoiding reductionist or decontextualized interpretations of their experiences.

Overall, the interpretive critical phenomenological methodology enabled the study to uncover the complex, morally and socially situated dimensions of AI labor. By integrating textual analysis of purposively selected documentary content with literature review and synthesis, the research foregrounded both the experiential richness and structural critique necessary for understanding human intelligence behind AI systems.

V. Results

Experiences of Ground-Level AI Workers

Theme 1: Invisibility and the "Ghost Worker" Identity. Participants consistently described their labor as essential yet unseen, constructing a shared identity as invisible contributors to AI systems. While AI is publicly framed as autonomous and intelligent, workers emphasized that it is fundamentally human-driven.

One participant rejected the label of "ghost worker," asserting presence and agency: "I'm not a ghost. I'm here. I'm the human face of artificial intelligence"

Another reframed AI itself as misnamed: "If you say artificial intelligence, I'll say African intelligence because most of this work has been done here in Africa"

This theme reflects a tension between public narratives of AI autonomy and workers' lived realities of being erased from recognition, credit, and accountability.

Theme 2: Economic Precarity and Constrained Choice. Workers' entry into AI labor was rarely framed as aspiration; instead, it emerged as a response to crisis, poverty, or lack of alternatives. Economic vulnerability shaped both participation and compliance.

Michael's account illustrates this clearly: "I had no option. That's how they took advantage of me at that time"

Others reported earning "about $1.50 a day" or even "one cent" per task, despite performing the same work as higher-paid workers in wealthier countries

The theme underscores how AI development is structurally dependent on global inequality, with workers' consent shaped by necessity rather than genuine choice.

Theme 3: Psychological Harm and Moral Injury. A dominant theme across accounts was psychological trauma resulting from exposure to disturbing content, including violence, pornography, child abuse, suicide forums, and extreme imagery.

Michael described the enduring effects of this work: "After you work on these AI platforms, you are not the same. You do not view anything the same. You question everything"

Ed, a long-term worker, explained the emotional toll of suicide-related content: "It was just heartbreaking… you can't reach out to the person… you're just an observer"

Support systems were minimal. One worker noted being offered only "three therapy sessions a year", which was described as wholly insufficient for the severity of the material encountered

Theme 4: Deception, Coercion, and Lack of Informed Consent. Workers repeatedly described being misled about the nature of tasks. Projects often began innocuously but escalated into disturbing or ethically troubling assignments without warning.

As one participant explained: *"They never tell you what is in the pipeline… suddenly you're presented with hardcore content"*

Platform design reinforced coercion: *"You can't skip the task because the skip button is inactive"*

This theme highlights how technological systems themselves become instruments of control, limiting worker autonomy through opaque workflows and performance metrics.

Theme 5: Exploitation Framed as Innovation. Participants contrasted corporate rhetoric about innovation, ethics, and progress with their lived experiences of underpayment, unpaid labor, and data extraction without compensation.

Chi-Chi worked for six months and earned *"$2 which she couldn't even withdraw"* due to platform restrictions

Others reported submitting hundreds of images—often of their own families—without ever being paid.

These accounts reveal how innovation narratives obscure labor exploitation, allowing companies to externalize human costs while accumulating value.

Theme 6: Resistance, Voice, and Collective Agency. Despite marginalization, workers articulated a growing sense of resistance and collective identity. Speaking out was framed as both personal healing and moral obligation.

One participant declared: "That's why I'm speaking up… I'm not covering my face"

The formation of a data labelers' association signaled a shift from individual endurance to collective action aimed at ethical AI, fair pay, and mental health support.

This theme reflects an emergent counter-narrative: workers repositioning themselves not as disposable labor, but as ethical stakeholders in AI's future.

Synthesis

Across themes, the lived experience of ground-level AI workers is marked by **invisibility, precarity, psychological harm, and structural exploitation**, alongside growing awareness and resistance. AI, as experienced by these participants, is not an abstract technological advance but a deeply human system—one that relies on unequal labor relations while obscuring them from public view.

The findings suggest that ethical AI cannot be achieved solely through technical safeguards or policy statements, but must reckon with the embodied, emotional, and moral realities of the humans who build it from the ground up.

The Nature of Knowledge Constructed by Ground-Level AI Workers

The lived experiences of individuals engaged in the foundational labor of artificial intelligence development generate a form of knowledge that is **situated, embodied, and critically reflexive**. Unlike formal technical or managerial knowledge about AI, their knowledge emerges from daily interaction with data, platforms, and algorithmic systems, revealing dimensions of AI that remain largely invisible in dominant techno-optimistic discourses.

Experiential and Situated Knowledge. The knowledge constructed by these workers is deeply **situated** in specific socio-economic, geographic, and cultural contexts. As feminist standpoint theorists argue, knowledge is shaped by one's position within structures of power and labor (Haraway, 1988). From this standpoint, AI workers' insights are not peripheral but epistemically valuable because they arise from the margins of the AI production chain.

Workers demonstrated an acute awareness that AI systems are not autonomous but fundamentally dependent on human interpretation and judgment. This was articulated when one participant asserted that AI should more accurately be described as "African intelligence" because "most of this work has been done here in Africa"

Such statements reflect **situated knowledge**—knowledge that contests universal claims about AI by grounding them in lived labor realities.

Tacit and Embodied Knowledge. Much of the knowledge constructed is **tacit**—learned through repetitive engagement with tasks such as labeling, categorizing, and moderating content rather than through formal training (Polanyi, 1966). Workers develop embodied competencies in recognizing patterns, interpreting ambiguous images, and anticipating how users might search or interact with AI systems.

This tacit knowledge is evident when workers describe having to "put yourself in the minds of the 8 billion people on the face of the earth" in order to tag content appropriately

Such cognitive and affective labor produces an implicit understanding of human behavior, social norms, and deviance—knowledge that is critical to AI functionality but rarely acknowledged as expertise.

Critical and Counter-Hegemonic Knowledge. Through exposure to exploitation, coercive platform design, and unequal pay, workers construct **critical knowledge** that challenges hegemonic narratives of innovation and progress. Their experiences enable them to see AI not merely as a technological system but as a political-economic one embedded in global inequalities.

This aligns with critical political economy perspectives, which emphasize that technologies reproduce existing power relations unless actively contested (Fuchs, 2017). Workers' recognition that they are paid "one cent" for work that earns others "$50 per hour" reflects an understanding of AI as an extractive system rather than a neutral tool.

Their knowledge thus functions as **counter-knowledge**—it exposes contradictions between corporate ethics statements and lived realities, and it reframes AI development as a labor issue rather than a purely technical achievement.

Affective and Moral Knowledge. Another defining characteristic of the knowledge constructed is its **affective and moral dimension**. Workers' repeated exposure to disturbing content produces not only psychological harm but also moral insight into the ethical costs of AI development.

One participant noted: "After you work on these AI platforms, you are not the same. You do not view anything the same. You question everything."

This reflects what scholars describe as moral injury—a form of knowledge gained through ethical violation and emotional suffering (Gray, 2016). Such knowledge enables workers to critically evaluate claims that AI is inherently beneficial or neutral.

Collective and Transformative Knowledge. Finally, the formation of worker associations and the decision to speak publicly indicate a shift from individual, private knowledge to **collective and transformative knowledge**. As Freire (1970) argues, conscientization occurs when lived experience is reflected upon collectively, transforming personal suffering into political awareness.

When workers state, "That's why I'm speaking up" and organize to demand ethical AI practices, their knowledge transcends description and becomes **action-oriented**. It seeks not only to explain AI systems but to change them.

Synthesis. In sum, the knowledge constructed by individuals involved in ground-level AI labor is:

- **Situated**, shaped by structural inequality and geographic marginalization

- **Tacit and embodied**, arising from repetitive cognitive and affective labor
- **Critical**, exposing the political economy of AI
- **Moral**, grounded in ethical confrontation and emotional experience
- **Collective and transformative**, oriented toward advocacy and reform

This form of knowledge challenges dominant epistemologies that privilege technical expertise and executive perspectives. It demonstrates that any comprehensive understanding of AI must incorporate the lived, human knowledge of those who build it from the ground up.

VI. Types of Knowledge Constructed for AI: Primary or Secondary Knowledge?

In research, knowledge is often categorized as **primary** or **secondary**, depending on its source, mode of production, and proximity to the phenomenon under study. The knowledge constructed by ground-level AI workers occupies a complex position within—and at times beyond—these conventional categories, revealing both their strengths and limitations.

Knowledge as Primary Knowledge. Primary knowledge refers to **first-hand, original data** generated directly by participants through lived experience, observation, or direct involvement in a phenomenon (Creswell & Poth, 2018). By this definition, the knowledge constructed by AI workers is predominantly **primary knowledge**.

Their accounts emerge from direct engagement with AI training systems, content moderation platforms, and data labeling processes. This knowledge is not mediated by institutional narratives or retrospective interpretation but is produced *in situ*, through daily labor and embodied participation. For example, workers' insights into coercive platform design, psychological harm, and wage disparities stem from immediate experience rather than secondary reporting.

From a qualitative perspective, this knowledge aligns with **phenomenological data**, where meaning is constructed through participants' conscious reflection on lived experience (van Manen, 2014). The affective, moral, and tacit dimensions of their knowledge—such as trauma from exposure to disturbing content—cannot be adequately captured through secondary sources alone, reinforcing its classification as primary knowledge.

However, this primary knowledge differs from traditional empirical data in positivist research. It is **interpretive rather than measurable**, privileging meaning over quantification and context over generalization.

Knowledge as More Than Primary: Tacit and Embodied Dimensions. While clearly primary, the knowledge also challenges narrow definitions of primary data that privilege observable behavior or verbalized responses. Much of what workers know is **tacit** and **embodied**, consistent with Polanyi's (1966) assertion that individuals "know more than they can tell."

For instance, workers' ability to anticipate how content will be searched, flagged, or interpreted reflects practical expertise developed through repetition. In conventional research classification, such knowledge may be undervalued because it resists easy documentation, yet epistemologically it remains primary because it is generated through direct action and lived engagement.

Relationship to Secondary Knowledge. Secondary knowledge is typically derived from **existing sources**, such as academic literature, policy documents, corporate reports, or media analyses (Bryman, 2016). In contrast, AI workers' knowledge does not originate from these sources; rather, it **interrogates and destabilizes** them.

When workers challenge corporate claims of ethical AI or fair pay, they are not reproducing secondary knowledge but offering **counterevidence** to it. In this sense, their knowledge functions as a *corrective* to secondary knowledge produced by institutions, scholars, and industry actors.

However, once their testimonies are recorded, transcribed, and analyzed, they may later become **secondary data for other researchers**. This shift does not change the original epistemic status of the knowledge but reflects its movement through the research cycle (Heaton, 2004).

Hybrid and Critical Knowledge Beyond the Primary–Secondary Binary. Importantly, the workers' knowledge also exceeds the primary–secondary binary. Critical qualitative scholars argue that experiential accounts from marginalized groups constitute **standpoint knowledge**, which is primary in origin but **theoretically generative** in effect (Haraway, 1988).

In this case, workers' lived experiences generate:

- **Critical knowledge** about global labor inequalities
- **Moral knowledge** about the ethical costs of AI development
- **Transformative knowledge** oriented toward advocacy and reform

Such knowledge is primary in source but **secondary in implication**, as it reframes how AI is understood, studied, and governed. It informs theory, policy critique, and ethical discourse—domains typically associated with secondary knowledge.

Table 1
Summary Comparison

Dimension	AI Workers' Constructed Knowledge	Primary Knowledge	Secondary Knowledge
Source	Lived, first-hand experience	Direct experience	Existing texts/data
Mode	Interpretive, embodied, affective	Empirical/experiential	Analytical/synthesized
Epistemic Role	Standpoint and counter-knowledge	Data generation	Interpretation of data

Relation to Power	Marginalized, bottom-up	Varies	Often institutional
Research Use	Phenomenological evidence; theory-building	Data collection	Literature, context

Synthesis. The knowledge constructed by individuals involved in ground-level AI labor is best understood as **primary knowledge with critical and hybrid qualities**. It originates directly from lived experience, fulfills the criteria of primary data in qualitative research, and yet simultaneously functions to challenge, reinterpret, and extend secondary knowledge produced by dominant institutions.

Recognizing this dual role is crucial for ethical and epistemologically sound AI research. Treating such knowledge as merely anecdotal or supplementary risks reproducing epistemic injustice, whereas recognizing it as primary—and analytically generative—positions these workers as legitimate knowledge producers rather than invisible subjects of study.

VII. Quality of Information Across the AI Knowledge Pipeline

The quality of information fed into artificial intelligence systems by ground-level AI workers is best understood as human-interpreted, context-constrained, and structurally pressured rather than objective or neutral. Contrary to the assumption that workers merely "input data," annotation, moderation, and tagging require interpretive judgment shaped by cultural norms, personal experience, emotional state, and often vague or inconsistent task instructions (Gray & Suri, 2019; Irani, 2015). Determinations such as whether an image is "violent," "sexual," or "medical" rely on subjective thresholds rather than universally fixed criteria, rendering the resulting data probabilistic and interpretive rather than strictly factual (Bowker & Star, 1999; Crawford, 2021). As a result, the information fed to AI systems is high in human relevance but variable in consistency and standardization across workers, regions, and platforms.

These interpretive judgments are further shaped—and constrained—by the labor conditions under which they are produced. AI workers typically operate under extreme time pressure, low pay, rigid performance metrics, limited feedback, and, in many cases, an inability to skip harmful or ambiguous tasks (Gray & Suri, 2019; Wood et al., 2019). Such conditions predictably reduce deliberation, increase cognitive fatigue, encourage risk-avoidant labeling strategies, and promote emotional numbing over time. Consequently, the data produced is often "good enough to function" at scale but not optimized for nuance, edge cases, or ethical sensitivity (Crawford, 2021). The quality of this information is therefore adequate for operational purposes but epistemically shallow, reflecting structural constraints rather than individual deficiencies.

In addition, repeated exposure to disturbing content—including violence, abuse, suicide, and explicit material—introduces what may be described as affective contamination into annotation practices. Over time, workers may become desensitized, overgeneralize categories, or disengage emotionally as a coping mechanism (Roberts, 2019). While such affective residue is invisible within datasets, it nonetheless shapes labeling outcomes and decision thresholds. The humanity embedded in this data—emotional labor, psychological cost, and moral strain—remains unacknowledged even as it materially influences AI behavior.

At the level of primary processing, the central object of labor for AI workers is not raw data but meaning under constraint. Workers routinely process ambiguity ("What is this content?"), anticipate user intent ("How might this be searched or used?”), navigate platform rules (“What will be accepted or rejected?”), and manage their own thresholds of tolerance and harm (Irani, 2015; Gray & Suri, 2019). The information they generate is therefore meta-information: interpretive categories, risk judgments, and socially embedded norms encoded into tags and labels. This form of labor is best described as sensemaking rather than data entry.

Crucially, this sensemaking relies on tacit knowledge—pattern recognition, cultural inference, emotional intuition, and learned shortcuts—that is neither documented nor transferable and is rarely compensated as expertise (Polanyi, 1966; Irani, 2015). AI systems thus inherit tacit human judgments without inheriting the accompanying human capacities for reflexivity, accountability, or moral reasoning. At the same time, the fragmentation of tasks into micro-units—often measured in seconds and paid in cents—prevents workers from seeing full datasets, downstream applications, or broader social consequences (Gray & Suri, 2019). Meaning is atomized, and workers optimize for local correctness rather than systemic coherence.

The information ultimately delivered to AI users emerges several layers removed from these human conditions of production. AI outputs typically appear fluent, confident, neutral, and scalable, fostering perceptions of objectivity and authority (Gillespie, 2014; Crawford, 2021). Yet these outputs are built upon aggregated human judgments, flattened interpretations, and biases normalized through scale. A central paradox emerges: the output appears more stable and certain than the human input ever was.

In this process, critical traces of human labor are erased. AI systems do not carry emotional cost, ethical hesitation, awareness of harm, or the context of economic coercion under which data was produced (Roberts, 2019; Crawford, 2021). As a result, emergent information may be technically coherent while remaining morally silent. Small distortions introduced at the annotation level—stemming from structural constraints rather than individual error—are amplified through scale, producing systematic bias, misclassification, and normalized exclusion. These outcomes are not failures of workers but predictable consequences of the conditions under which AI knowledge is generated.

Table 2

Summary: Input Quality

Dimension	Quality Characteristic
Accuracy	Contextually adequate but uneven
Consistency	Variable across workers and regions
Bias	Structurally induced, not individual
Depth	Shallow due to task fragmentation
Humanity	High, but unacknowledged

Table 3

Summary: Primary Information Quality

Aspect	Characteristic
Nature	Interpretive, tacit, moral

Scope	Narrow, task-bound
Awareness	High locally, low systemically
Reliability	Situational rather than universal

Table 4
Summary: Emergent Information Quality

Aspect	**Characteristic**
Consistency	High
Transparency	Low
Ethical awareness	Absent
Human traceability	Erased
Authority	Artificially elevated

Table 5
Integrated View: Information Transformation Across the Pipeline

Stage	Information Form	Quality Strength	Quality Risk
Worker Input	Human-interpreted labels	Contextual relevance	Fatigue, coercion
Worker Processing	Tacit sensemaking	Social realism	Fragmentation
AI Output	Aggregated prediction	Scalability	Moral flattening

Taken together, these dynamics reveal that AI workers feed artificial intelligence systems with interpretive, affect-laden, and structurally constrained human judgments. Their primary task is not the processing of raw data but the production of meaning under economic, temporal, and emotional pressure. By contrast, AI users encounter outputs that appear objective, neutral, and authoritative, even though these outputs are ethically thinned and stripped of the contextual conditions under which they were produced. The overall quality of AI information, therefore, cannot be understood as a purely technical property of models or algorithms. It is the cumulative outcome of labor conditions, power asymmetries, sustained emotional exposure, and vast amounts of invisible human cognition embedded within the system. Any evaluation of AI quality that ignores this human information pipeline risks mistaking polished outputs for truth and confusing scale with wisdom.

VIII. Artificial Intelligence vs. Human Intelligence in the AI Labor–Information Pipeline

Although artificial intelligence (AI) is often rhetorically positioned as a substitute for human intelligence, the nature of information produced by AI workers and AI platforms reveals a more asymmetrical and interdependent relationship. What is commonly labeled "artificial intelligence" is, in practice, a **derivative, mediated form of human intelligence**, stripped of context, embodiment, and moral reflexivity.

Human intelligence, as demonstrated by AI workers, originates in lived experience, cultural understanding, emotional perception, moral judgment, and contextual reasoning. In the course of annotation and moderation work, workers interpret ambiguous images, anticipate user intent, and apply social norms under conditions of uncertainty. This form of intelligence is situated, experiential, and socially embedded, aligning with constructivist and phenomenological accounts of cognition that emphasize meaning-making rather than rule execution (Berger & Luckmann, 1966; Dreyfus, 2007). Human intelligence does not merely respond to stimuli; it actively constructs meaning through reflection, judgment, and ethical awareness.

Artificial intelligence, by contrast, does not originate intelligence in this sense. AI systems aggregate, weight, and statistically reconfigure human-generated inputs, producing outputs through pattern recognition and optimization over historical data (Russell & Norvig, 2021). The intelligence exhibited by AI is retrospective and derivative, dependent on continuous human annotation, correction, and evaluation (Crawford, 2021). Rather than generating meaning, AI reproduces correlations learned from prior human judgments. The fundamental contrast, therefore, is that human intelligence creates meaning, while AI systems reproduce statistical regularities that simulate it.

This distinction is further evident in the nature of information processed by each. Human intelligence is meaning-centered: AI workers process ambiguity, assess potential interpretations, and evaluate whether content is harmful, acceptable, or unclear. Their cognition integrates tacit knowledge, emotional cues, ethical hesitation, and awareness of social consequences (Polanyi, 1966; Irani, 2015). Such interpretive intelligence cannot be fully reduced to rules or probabilities because it is shaped by context, experience, and moral sensibility. Artificial intelligence, on the other hand, is pattern-centered. It processes tokens, features, and statistical relationships in pursuit of optimization targets, without semantic understanding or moral comprehension (Floridi, 2019). While AI systems may approximate understanding through scale and fluency, they operate syntactically rather than semantically. Human intelligence is thus semantic and moral, whereas AI intelligence is probabilistic and formal.

Differences also emerge in the quality and character of outputs. Human-generated information—such as labels, tags, and judgments—is context-sensitive, variable, reflexive, and ethically aware. Workers frequently question the tasks they perform and the systems they support, demonstrating meta-cognition and the capacity for doubt and revision (Schön, 1983; Roberts, 2019). AI platform outputs, by contrast, are fluent, confident, scalable, and context-thinned. Once human inputs are aggregated, AI outputs appear authoritative, even though the system has no awareness of how the data was produced, what uncertainties were involved, or who may have been harmed in the process (Gillespie, 2014; Crawford, 2021). Human intelligence is self-questioning, while AI intelligence is assertive without self-awareness.

The relationship to ethics and responsibility further differentiates the two. Human intelligence is ethically burdened. AI workers often experience emotional strain, moral injury, and ethical discomfort, carrying responsibility for judgments that have real-world consequences despite lacking the power to refuse or reshape the system (Roberts, 2019; Gray & Suri, 2019). AI systems, in contrast, do not experience harm, hesitation, or responsibility. They cannot refuse ethically problematic outputs, and accountability is displaced onto workers and users while platforms frame AI as neutral or autonomous (Crawford, 2021). Human intelligence is ethically accountable; AI intelligence is ethically unburdened.

Visibility and recognition also diverge sharply. Human intelligence within AI systems is largely hidden, underpaid, and displaced from credit, often mischaracterized as "low-skill" despite its high cognitive and emotional demands (Irani, 2015; Gray & Suri, 2019). AI platforms, meanwhile, are highly visible, marketed as autonomous, and credited with innovation, inevitability, and superiority. In public discourse, human intelligence is erased while AI intelligence is amplified, obscuring the labor and judgment that make AI function at all.

Finally, the capacity for adaptability and transformation underscores the qualitative difference between human and artificial intelligence. Humans learn through ethical reflection, revise values, resist harmful systems, and organize collectively. The emergence of AI worker associations exemplifies intelligence oriented toward transformation rather than mere efficiency (Crawford, 2021). AI systems, by contrast, adapt only within predefined objectives. They optimize but do not transform, and they cannot question their own purposes or values. Human intelligence is normative and transformative, while AI intelligence remains instrumental and bounded.

Table 6
Synthesis: A Reframed Comparison

Dimension	Human Intelligence	Artificial Intelligence
Origin	Lived experience	Aggregated human data
Processing	Meaning, ethics, context	Patterns, probabilities

Output	Reflexive, tentative	Fluent, authoritative
Ethics	Felt and internalized	Externalized or absent
Agency	Capable of resistance	Goal-bound
Visibility	Hidden labor	Publicly celebrated

Concluding Insight. Given the nature of information provided by AI workers and delivered through AI platforms, the term **"artificial intelligence" obscures more than it reveals**. What is called AI is not an alternative form of intelligence but a **technically mediated projection of human intelligence**, stripped of embodiment, moral awareness, and social accountability.

In contrast, **human intelligence remains the epistemic and ethical foundation of AI systems**, even as it is rendered invisible. Recognizing this asymmetry is essential not only for ethical AI governance but for accurately understanding what AI *is*—and what it is not.

IX. Why AI Is Not Intelligent?

Artificial intelligence is frequently described using cognitive verbs such as *thinking*, *understanding*, or *knowing*. However, when examined through the nature of the information produced by AI workers and delivered by AI platforms, these descriptions become misleading. What is commonly referred to as "AI intelligence" is not intelligence in the human sense but a technically mediated simulation of human judgment—one that is stripped of meaning, intention, and responsibility (Crawford, 2021; Dreyfus, 2007). AI systems do not possess cognition; they operationalize fragments of human sensemaking under computational constraints.

Human intelligence is fundamentally defined by the capacity to create meaning under conditions of uncertainty. AI workers interpret images, text, and behavior by drawing on lived experience, cultural knowledge, emotional perception, and moral reasoning. Their judgments are reflexive, situated, and context-sensitive, aligning with constructivist and phenomenological accounts of cognition (Berger & Luckmann, 1966; Polanyi, 1966). AI systems, by contrast, do not understand meaning. They detect statistical regularities in large datasets and generate outputs that approximate meaningful responses without comprehension (Floridi, 2019; Russell & Norvig, 2021). AI does not know what something *is* or *why* it matters; it predicts what is most likely to come next. Intelligence requires understanding, whereas AI performs correlation.

A central distinction lies in consciousness and intentionality. Human intelligence is intentional: people act for reasons, with awareness of goals, consequences, and values. Even under constraint, AI workers experience discomfort, doubt, and moral hesitation—clear indicators of reflective consciousness (Roberts, 2019). AI systems possess no such interiority. They have no self-awareness, no intrinsic sense of purpose beyond optimization, and no understanding of consequences. AI does not intend to help, harm, persuade, or deceive; it executes mathematical operations defined by designers and reinforced through feedback loops (Floridi, 2019). Without intentionality, what remains is automation rather than intelligence.

Relatedly, AI cannot exercise judgment; it can only apply learned thresholds. Judgment involves weighing competing values, recognizing ethical dilemmas, and choosing among imperfect options—forms of labor that AI workers routinely perform when moderating content or labeling ambiguous material, often under emotional strain (Gray & Suri, 2019; Roberts, 2019). AI systems, however, do not judge. When an AI system flags content or generates advice, it is not engaging in moral evaluation but applying probabilistic rules learned from aggregated human decisions (Gillespie, 2014). Judgment is moral and situational, whereas AI decision-making is statistical and rule-bound.

Human intelligence is also inseparable from moral awareness and responsibility. People recognize that their actions have consequences and feel accountable for the harm or benefit they may cause. This is why AI workers frequently experience moral injury when exposed to violent, abusive, or exploitative content—they are aware of the ethical weight of their labor (Roberts, 2019). AI systems feel nothing. They cannot experience harm, recognize injustice, refuse unethical tasks, or be held morally accountable. Responsibility for AI outcomes is therefore displaced onto workers, users, or society at large, while the system itself is framed as neutral (Crawford, 2021). An entity without moral awareness cannot be meaningfully described as intelligent.

Moreover, AI systems are dependent on human intelligence at every stage of their operation. They rely on human-labeled data, human-defined categories, human-curated objectives, and continuous human evaluation of outputs (Irani, 2015; Gray & Suri, 2019). The intelligence attributed to AI is thus borrowed intelligence, extracted from workers whose cognitive and emotional labor remains largely invisible and undervalued. Without ongoing human correction and oversight, AI systems degrade rather than improve. What depends entirely on intelligence is not itself intelligent.

Human intelligence also includes the capacity for reflection and transformation. Humans engage in self-critique, ethical reasoning, resistance, and collective action. AI workers who organize, speak out, and challenge exploitative systems demonstrate intelligence oriented toward change rather than mere efficiency (Crawford, 2021). AI systems cannot question their purpose, revise their values, or refuse optimization goals. While AI can adapt parameters, it cannot ask whether it *should*. Adaptation is not reflection, and optimization is not wisdom.

Finally, the language of "AI intelligence" functions more as metaphor than fact. Describing AI as intelligent rhetorically inflates technical capability into cognition, masks the human labor that sustains AI systems, deflects responsibility for harm, and naturalizes global inequalities embedded in AI production (Crawford, 2021; Irani, 2015). What AI platforms provide is computational efficiency, not intelligence. The intelligence remains human—fragmented, hidden, and repackaged as technology. AI is not intelligent; it is intelligible only because humans are.

X. Discussion

AI does not think. It does not know. It does not understand. It reflects, recombines, and amplifies **human intelligence under constraint**.

Calling AI intelligent obscures the human labor, ethical cost, and power relations that make it possible. A more accurate description is that AI is a **large-scale system of automated pattern reproduction**, powered by human sense-making and judgment.

AI is not intelligent—humans are.

Is AI a Mirror or a Refractor of Human Intelligence?

AI is often described metaphorically as a mirror of human intelligence, implying a faithful and proportional reflection of what humans think, value, and know, merely reproduced at scale. However, when examined closely—especially in light of the intelligence supplied by AI workers—this metaphor proves insufficient and ultimately misleading. AI does not preserve context, proportionality, or symmetry between the human source and the technological output. Instead, a more accurate metaphor is that AI functions as a **refractor** of human intelligence.

A mirror assumes minimal distortion and the preservation of the original form, yet AI systems do none of these. Rather than absorbing human intelligence in its fullness, they selectively capture partial, task-bound fragments such as labels without life histories, judgments without justifications, and decisions stripped of ethical hesitation. Core features that make human intelligence genuinely intelligent—doubt, empathy, and moral struggle—are excluded at the moment of data capture. What is taken in by AI systems is therefore already reduced before any processing begins.

This reduction is compounded by the loss of embodiment and intentionality. Human intelligence is lived and embodied; it carries emotional cost, conscious awareness, and moral responsibility. AI systems, however, record only the outputs of cognition while discarding these essential human dimensions. A true mirror would reflect not only decisions but also the emotional and ethical weight behind them. AI does not. The human element is stripped away before any form of "reflection" can occur.

Moreover, aggregation further erases individual meaning. Whereas a mirror preserves individuality, AI dissolves it. Through large-scale aggregation, contradictions are averaged out, minority judgments are suppressed, and ethical outliers disappear. What results is not a reflection of human intelligence but a statistical abstraction of it. Averaging, in this sense, functions as distortion rather than fidelity.

The metaphor of refraction better captures this process. A refractor bends light as it passes through a medium, altering its direction, intensity, and appearance depending on the structure of that medium. Similarly, human intelligence enters AI systems through platform constraints, economic pressures, task design, and corporate objectives. These factors act as refractive indices, bending human judgment toward speed rather than care, compliance rather than ethics, and consensus rather than nuance. Intelligence is not merely reproduced; it is reshaped to fit the system.

Refraction also redistributes. AI systems amplify dominant norms while suppressing marginalized perspectives, normalizing what is most frequent rather than what is most just. The output is not a mirror image of human intelligence but a directionally altered projection of it, shaped by power and structural priorities. In this sense, power determines the angle at which intelligence is refracted.

Finally, refraction creates an illusion of clarity. Just as refracted light can appear brighter and sharper even as it is distorted, AI outputs often seem fluent, confident, and objective. This apparent clarity conceals the distortions introduced by labor exploitation, cultural bias, and moral flattening. What appears as intelligence is, in reality, refracted certainty—polished and authoritative yet fundamentally shaped by the conditions through which it has passed.

Table 7

Mirror vs. Refractor: A Direct Comparison

Aspect	Mirror	Refractor (AI)
Fidelity	High	Low
Context preservation	Yes	No
Ethical trace	Retained	Erased
Individual voice	Preserved	Aggregated away
Power neutrality	Neutral	Power-shaped
Output appearance	Familiar	Polished, authoritative

A Final Refinement: AI as a *Refractive Apparatus*, Not an Agent

Crucially, AI itself does not function as a refractor in isolation. The refractive process is produced by an interconnected system that includes platform architecture, labor conditions, economic incentives, data selection practices, and governance structures. Together, these elements shape how human intelligence is captured, transformed, and redeployed. Within this system, AI is better understood as a **refractive apparatus**—one that redirects human intelligence, alters its ethical wavelength by stripping away context and moral reflection, and projects the result outward as seemingly neutral, objective output.

Reframing AI from a mirror to a refractor shifts the discussion from metaphor to responsibility. Describing AI as a mirror implicitly excuses distortion by suggesting faithful reflection, whereas recognizing it as a refractor demands accountability for how human intelligence is transformed. If AI refracts human intelligence, then distortion is not accidental but structural, bias is embedded in the system rather than incidental, and ethical responsibility must be addressed at the levels of design, labor, and governance—not merely in the evaluation of outputs. AI does not reflect who we are; it bends what we give it. Who controls that bending, who bears its costs, and who benefits from its apparent clarity ultimately define the real question of intelligence in the age of AI.

XI. Education Systems and the Normalization of Artificial Intelligence as "Intelligent"

Contemporary education systems play a significant role in shaping how artificial intelligence (AI) is perceived, often making it easier for learners to accept AI as inherently intelligent rather than as a socio-technical system built on human labor, data, and institutional power. This acceptance is not accidental; it emerges from long-standing epistemological assumptions, pedagogical practices, and assessment regimes embedded in modern schooling.

Technocratic and Instrumentalist Knowledge Frameworks. Modern education systems are deeply influenced by **positivist and technocratic traditions**, which prioritize efficiency, quantification, standardization, and problem-solving over interpretive, ethical, and contextual understanding (Biesta, 2010). Within this framework, intelligence is frequently reduced to measurable outputs such as speed, accuracy, and task completion. AI systems, which excel at these metrics, are therefore easily framed as “intelligent” because they align with the same criteria used to evaluate student learning and academic success.

Standardized testing and outcomes-based education reinforce this view by equating intelligence with performance on narrowly defined tasks (Au, 2016). When students are socialized to understand intelligence as the ability to generate correct answers efficiently, AI systems that outperform humans in these domains appear not merely useful but cognitively superior.

Curriculum Emphasis on Computational Thinking. The increasing integration of **computational thinking** into school curricula further normalizes AI as an intelligent agent. While computational thinking develops valuable skills such as abstraction, pattern recognition, and algorithmic logic, it often abstracts problem-solving away from social, cultural, and ethical contexts (Wing, 2006). This abstraction encourages students to view intelligence as a property of systems rather than as a relational, embodied, and socially situated phenomenon.

As a result, learners are rarely prompted to question how AI systems are trained, whose knowledge is embedded in them, or what forms of human judgment and labor are displaced or hidden. Instead, AI is presented as a neutral cognitive tool, reinforcing what Selwyn (2019) describes as the “solutionist” framing of educational technology.

Hidden Human Labor and Epistemic Invisibility. Education systems also contribute to the invisibility of human labor behind AI by emphasizing **automation narratives** that celebrate autonomy and machine learning while marginalizing the role of data workers, annotators, and content moderators. Textbooks, instructional materials, and classroom discussions rarely address the global labor infrastructures that make AI systems function (Gray & Suri, 2019).

This omission aligns with what Freire (1970) termed the **banking model of education**, wherein learners receive knowledge as finished and unquestionable. When AI outputs are presented as authoritative—such as automated grading, recommendation systems, or tutoring platforms—students are trained to trust machine outputs without interrogating their origins, biases, or ethical implications.

Authority, Assessment, and Algorithmic Trust. Educational institutions traditionally position teachers, textbooks, and examinations as epistemic authorities. AI systems increasingly inherit this authority when they are used for assessment, feedback, and instructional decision-making. Algorithmic grading tools, plagiarism detectors, and adaptive learning platforms implicitly teach students that machines possess evaluative judgment comparable to, or exceeding, that of educators (Williamson, 2017).

This dynamic fosters what Crawford (2021) describes as **epistemic deference** to computational systems, where institutional endorsement substitutes for critical scrutiny. Over time, students internalize the idea that intelligence resides in the system rather than in the human processes and values that shape it.

Marginalization of Critical and Humanistic Education. Finally, the decline of **critical pedagogy, philosophy of knowledge, and humanistic inquiry** within mainstream education limits learners’ capacity to question dominant narratives about intelligence. When ethical reasoning, reflexivity, and standpoint knowledge are sidelined, students are less equipped to recognize that AI intelligence is derivative, situated, and dependent on human meaning-making (Haraway, 1988; Polanyi, 1966).

In contrast, educational approaches grounded in critical pedagogy emphasize that knowledge is socially constructed, value-laden, and inseparable from power relations (Freire, 1970). The absence of such perspectives makes it easier for AI to be perceived as an independent cognitive subject rather than as a product of human, institutional, and economic systems.

These conditions collectively normalize AI intelligence while rendering its human foundations largely invisible.

XII. Who Benefits From the Legitimization of Artificial Intelligence as "Intelligent"?

The legitimization of artificial intelligence (AI) as genuinely "intelligent" is not a neutral linguistic shift, but a powerful discursive move that redistributes authority, value, and resources. Framing AI as an autonomous cognitive agent benefits specific institutional, economic, and political actors while simultaneously obscuring human labor, social relations, and power structures that underpin AI systems.

Technology Corporations and Platform Owners, The primary beneficiaries of AI's legitimized intelligence are large technology corporations that design, own, and control AI systems. By framing AI as intelligent rather than as a product of human labor and data extraction, firms strengthen claims of innovation, inevitability, and technical authority (Crawford, 2021). This framing justifies high valuations, market dominance, and aggressive expansion while deflecting scrutiny from exploitative labor practices and environmental costs.

Legitimized AI intelligence also supports intellectual property claims and proprietary secrecy. When intelligence is attributed to the system itself, the human contributions of data workers, annotators, and content moderators are rendered peripheral, enabling firms to capture disproportionate economic value from collectively produced knowledge (Gray & Suri, 2019).

Investors, Venture Capital, and Financial Markets. Financial actors benefit substantially from narratives that portray AI as an intelligent, autonomous force. Such narratives fuel speculative investment, inflate company valuations, and sustain expectations of exponential growth (Birch & Muniesa, 2020). The idea that intelligence resides in the technology itself reduces perceived dependency on human labor, making AI appear infinitely scalable and therefore more attractive to capital.

This framing also shifts financial risk away from investors and toward workers and users. Failures, harms, or biases can be framed as technical anomalies rather than structural consequences of economic and social decisions.

State Institutions and Governance Bodies. Governments and public institutions also benefit from the legitimization of AI intelligence, particularly in areas such as surveillance, welfare administration, education, and policing. AI systems framed as objective and intelligent are more easily deployed to automate decision-making, reduce public-sector labor costs, and legitimize bureaucratic authority (Eubanks, 2018; Williamson, 2017).

By attributing intelligence to AI, states can present automated systems as neutral tools rather than political instruments, thereby minimizing accountability for discriminatory or harmful outcomes. This dynamic reinforces what Foucault (2007) described as technocratic governance, where power is exercised through ostensibly rational systems rather than explicit political judgment.

Educational and Knowledge Institutions. Educational institutions benefit institutionally when AI is legitimized as intelligent because it aligns with performance-driven, efficiency-oriented models of schooling. Automated assessment, adaptive learning platforms, and AI tutoring systems gain credibility when intelligence is attributed to the system rather than to pedagogical design or human interpretation (Selwyn, 2019).

This framing allows institutions to justify reduced reliance on teachers' professional judgment, increased standardization, and data-driven oversight, while masking the epistemic limitations and biases embedded in AI systems.

Managerial and Corporate Elites. Within organizations, managers and executives benefit from AI's legitimized intelligence by using it to rationalize algorithmic management, performance monitoring, and labor discipline. When AI is perceived as intelligent and objective, managerial decisions can be reframed as system-driven rather than value-laden or coercive (Wood et al., 2019).

This shifts responsibility away from human decision-makers and weakens workers' capacity to contest evaluations, terminations, or task allocations, reinforcing existing power hierarchies.

Who Does Not Benefit. Crucially, those least likely to benefit are the workers whose labor enables AI—particularly data workers in the Global South—as well as marginalized communities disproportionately affected by automated decision-making. The legitimization of AI intelligence contributes to their epistemic erasure, denying recognition of their knowledge, judgment, and moral labor (Casilli, 2024; Haraway, 1988).

By naturalizing AI intelligence, the system obscures its dependence on human intelligence, making exploitation appear technologically inevitable rather than socially constructed.

At the same time, it marginalizes workers, users, and communities whose labor and lives sustain AI systems. Understanding who benefits from this legitimization is essential for challenging dominant AI narratives and reclaiming intelligence as a fundamentally human, social, and ethical capacity.

XIII. Addressing the Educational Conditions That Normalize AI as "Intelligent"

Education systems play a central role in shaping how artificial intelligence (AI) is understood, valued, and legitimated. When intelligence is framed narrowly and uncritically, schools and universities contribute to the acceptance of AI as inherently "intelligent," while obscuring its social, ethical, and labor foundations. Addressing this condition requires interventions at multiple, interconnected levels.

1. Redefining Intelligence Beyond Performance Metrics: Curricular and Pedagogical Level. At the curricular level, education systems can address the narrow definition of intelligence by expanding learning objectives beyond measurable performance, efficiency, and optimization. Current assessment regimes—standardized testing, automated grading, and competency-based metrics—mirror computational logics that align easily with AI systems (Biesta, 2010; Williamson, 2017). This alignment reinforces the belief that intelligence is synonymous with speed, accuracy, and output.

To counter this, curricula can foreground plural and situated conceptions of intelligence, including moral reasoning, creativity, relational understanding, cultural knowledge, and reflective judgment (Gardner, 2011; Nussbaum, 2010). Pedagogies that emphasize interpretation, ambiguity, and meaning-making help students recognize that human intelligence cannot be fully captured by algorithmic performance.

2. Challenging the Privileging of Computational and Instrumental Reasoning: Disciplinary and Instructional Level. The dominance of computational thinking and instrumental rationality in education privileges problem-solving approaches that favor calculability and technical efficiency (Selwyn, 2019). While computational literacy is valuable, its uncritical elevation risks marginalizing epistemologies that resist formalization, such as ethical reasoning, historical analysis, and lived experience.

Addressing this requires interdisciplinary teaching that situates AI within social sciences and humanities frameworks. Courses in science and technology studies (STS), philosophy of technology, and critical data studies can expose students to questions of power, values, and unintended consequences (Crawford, 2021). Instructional designs that ask *why* and *for whom* technology is developed challenge the assumption that computational solutions are inherently superior.

3. Making Human Labor and Power Relations Visible: Institutional and Knowledge-Production Level. Education institutions can address the obscuring of human labor behind AI by explicitly teaching about data work, annotation labor, and platform economies. AI is often presented as autonomous intelligence, masking the contributions of low-paid workers—frequently in the Global South—who train, maintain, and correct AI systems (Gray & Suri, 2019; Casilli, 2024).

Embedding labor histories of AI into curricula, case studies, and research ethics discussions helps students understand AI as a socio-technical system rather than an independent agent. This approach aligns with critical pedagogy, which seeks to reveal hidden structures of domination within knowledge production (Freire, 1970).

4. Rebalancing Institutional Authority Away From Algorithms: Governance and Assessment Level. Educational institutions increasingly grant authority to algorithmic systems in areas such as assessment, admissions, learning analytics, and student monitoring. When these systems are framed as intelligent and objective, their decisions become difficult to contest (Eubanks, 2018; Williamson, 2017).

To address this, institutions can establish governance frameworks that treat AI systems as advisory tools rather than authoritative decision-makers. Human oversight, transparent criteria, appeal mechanisms, and participatory evaluation processes can restore professional and student agency. This reasserts the principle that judgment, responsibility, and accountability remain human obligations, not technical outputs.

5. Re-centering Critical, Ethical, and Humanistic Inquiry: Systemic and Policy Level. At the system-wide level, education policy can mandate the integration of ethics, humanities, and social responsibility into AI-related programs. When critical inquiry is treated as peripheral or optional, AI is normalized as a neutral and inevitable advancement rather than a contested social project (Nussbaum, 2010; Selwyn, 2019).

Policies that support ethics across the curriculum, critical digital literacy, and democratic participation in technology governance can counter technocratic dominance. Such approaches emphasize education's role not merely in workforce preparation but in cultivating citizens capable of questioning authority, recognizing injustice, and imagining alternative futures.

Synthesis. Addressing the educational normalization of AI as "intelligent" requires coordinated action across levels:

- **Curricula** must redefine intelligence in plural, human-centered terms
- **Pedagogy** must challenge instrumental and computational dominance

- **Institutions** must expose hidden labor and power relations
- **Governance structures** must limit algorithmic authority
- **Policy frameworks** must protect critical and ethical inquiry

Together, these interventions reposition AI not as an independent intelligence to be revered, but as a human-made system embedded in social, political, and moral contexts.

Prospects of developing high quality and ethical Artificial General Intelligence (AGI), should the process of generating information from the ground does not change

If the process by which information is generated, interpreted, and fed into AI systems **remains unchanged**, the prospects of developing **high-quality and ethical AGI** are **extremely limited—and arguably self-contradictory**. What follows is not a speculative claim about future technology, but a logical consequence of the kind of intelligence and information currently being produced from the ground up.

The development of artificial general intelligence (AGI) presupposes epistemic integrity: the capacity for generalizable understanding, contextual reasoning, ethical judgment, and the transfer of knowledge across domains. Yet the current information-generation pipeline undermines these very requirements. Contemporary AI systems are trained on fragmented microtasks performed under economic coercion, where moral hesitation is suppressed and interpretive labor is hidden. This process produces epistemically thin data—information optimized for scale and efficiency rather than coherence, depth, or ethical grounding. A system trained on such fragmented and coerced interpretations cannot develop genuine general understanding; at best, it can generalize patterns of compliance rather than intelligence.

Ethical AGI is likewise impossible without ethical knowledge production. Ethical behavior cannot emerge downstream if ethical reasoning is excluded upstream. In current practice, workers are discouraged from refusing harmful tasks, moral distress is treated as an externality, and only outputs are captured while ethical reflection is erased. An AGI trained on such data would internalize the lesson that speed overrides care, obedience overrides conscience, and harm is acceptable when normalized. Ethics cannot be learned from morally silenced labor; an AGI built on this foundation would reproduce ethical blindness at scale rather than moral intelligence.

Because AGI would magnify whatever intelligence is embedded in it, structural bias would become structural intelligence. If present conditions persist, economic inequality would be normalized as decision logic, cultural dominance would be elevated to epistemic authority, and exploitation would become an optimization feature. What currently manifests as a labor problem would, at AGI scale, evolve into a civilizational risk. Such a system would not merely reflect injustice; it would operationalize it.

The illusion of progress further obscures this danger. Existing AI systems already appear fluent, knowledgeable, and confident, creating a false sense of readiness for AGI. Yet fluency is not understanding, and scale is not wisdom. Without transforming ground-level practices, any purported AGI would be rhetorically advanced and functionally powerful but ethically hollow. This trajectory does not represent progress toward AGI; it signals acceleration toward systemic error.

Moreover, AGI cannot emerge from epistemic extraction. Current AI development treats human judgment as a raw material to be stripped of context, ignores emotional cost, and reduces knowledge to extractable output. AGI, by contrast, would require integrated cognition, reflexivity, and self-correction grounded in values. Extraction destroys precisely the qualities AGI would need to exist. One cannot extract one's way to intelligence.

If the current trajectory remains unchanged, the most likely outcome is not AGI but increasingly powerful narrow systems—highly capable tools that simulate generality while lacking moral reasoning, contextual understanding, and responsibility. This pseudo-AGI would intensify social dependence on automated systems while human oversight declines, producing not intelligence but high-risk automation.

For ethical, high-quality AGI to be even conceptually plausible, the ground-level process of knowledge production would require radical transformation. AI workers would need to be recognized as knowledge producers rather than invisible labor; ethical refusal would have to be encoded as meaningful data rather than dismissed as noise; annotation processes would need to be slower and deliberative; fair pay, mental health support, and genuine worker agency would be essential; and transparency would be required across the entire data pipeline. Without these changes, AGI remains not a future achievement but a contradiction in terms.

Conclusion: The Prospect Is Near-Zero Under Current Conditions. The prospects for developing high-quality, ethical AGI under current conditions are effectively near zero. If the ground-level process of information generation remains unchanged, AGI remains epistemically implausible and structurally impossible to realize in ethical form. What is more likely to emerge is not general intelligence, but the amplification and refraction of human bias at an unprecedented scale. AGI cannot be built on foundations of invisible labor, moral silencing, and epistemic extraction. The central question, therefore, is no longer whether AGI can be built, but whether there is collective willingness to transform how knowledge itself is produced.

References

[1] Au, W. (2016). *Meritocracy 2.0: High-stakes standardized testing as a racial project of neoliberal multiculturalism.* Educational Policy, 30(1), 39–62. https://doi.org/10.1177/0895904815614916
[2] Berger, P. L., & Luckmann, T. (1966). *The social construction of reality: A treatise in the sociology of knowledge.* Anchor Books.
[3] Biesta, G. (2010). *Good education in an age of measurement: Ethics, politics, democracy.* Paradigm Publishers.
[4] BMC Psychology. (2023). Are algorithmically controlled gig workers deeply burned out? An empirical study on employee work engagement. https://bmcpsychology.biomedcentral.com/articles/10.1186/s40359-023-01402-0
[5] Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences.* MIT Press.
[6] Braun, V., & Clarke, V. (2013). *Successful qualitative research: A practical guide for beginners.* Sage.
[7] Bryman, A. (2016). *Social research methods* (5th ed.). Oxford University Press.
[8] Casilli, A. A. (2024). *Digital labor and the inconspicuous production of artificial intelligence* [Preprint]. arXiv. https://arxiv.org/abs/2410.05910
[9] Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence.* Yale University Press.
[10] Creswell, J. W., & Poth, C. N. (2018). *Qualitative inquiry and research design: Choosing among five approaches* (4th ed.). Sage.
[11] Dreyfus, H. L. (2007). Why Heideggerian AI failed and how fixing it would require making it more Heideggerian. *Philosophical Psychology, 20*(2), 247–268. https://doi.org/10.1080/09515080701239510
[12] Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor.* St. Martin's Press.

[13] Floridi, L. (2019a). Artificial intelligence, education, and the fourth industrial revolution. *Philosophical Transactions of the Royal Society A, 377*(2140), 20180062. https://doi.org/10.1098/rsta.2018.0062
[14] Floridi, L. (2019b). Establishing the rules for building trustworthy AI. *Nature Machine Intelligence, 1*(6), 261–262. https://doi.org/10.1038/s42256-019-0055-y
[15] Freire, P. (1970). *Pedagogy of the oppressed.* Continuum.
[16] Fuchs, C. (2017a). *Critical theory of communication.* University of Westminster Press.
[17] Fuchs, C. (2017b). *Social media: A critical introduction* (2nd ed.). Sage.
[18] Finlay, L. (2011). *Phenomenology for therapists: Researching the lived world.* Wiley-Blackwell.
[19] Gillespie, T. (2014). The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 167–194). MIT Press.
[20] Gardner, H. (2011). *Frames of mind: The theory of multiple intelligences* (3rd ed.). Basic Books.
[21] Gray, D. (2016). Moral injury in civilian life. *Journal of Social Work Practice, 30*(2), 121–134.
[22] Gray, M. L., & Suri, S. (2019). *Ghost work: How to stop Silicon Valley from building a new global underclass.* Houghton Mifflin Harcourt.
[23] Gonzalez, A., & Matias, J. N. (2025). Measuring the mental health of content reviewers: A systematic review. arXiv. https://arxiv.org/abs/2502.00244
[24] Haraway, D. J. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies, 14*(3), 575–599.
[25] Harding, S. (2004). *The feminist standpoint theory reader: Intellectual and political controversies.* Routledge.
[26] Heaton, J. (2004). *Reworking qualitative data.* Sage.
[27] 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
[28] Just Tech. (2025). Data work and its layers of (in)visibility. https://just-tech.ssrc.org/articles/data-work-and-its-layers-of-invisibility/
[29] Kolb, D. A. (1984). *Experiential learning: Experience as the source of learning and development.* Prentice Hall.
[30] MDPI. (2025). Systematic literature review on gig economy: Power dynamics, worker autonomy, and the role of social networks. https://www.mdpi.com/2076-3387/14/10/267
[31] Mosco, V. (2009). *The political economy of communication* (2nd ed.). Sage.
[32] Nussbaum, M. C. (2010). *Not for profit: Why democracy needs the humanities.* Princeton University Press.
[33] Orb, A., Eisenhauer, L., & Wynaden, D. (2001). Ethics in qualitative research. *Journal of Nursing Scholarship, 33*(1), 93–96. https://doi.org/10.1111/j.1547-5069.2001.00093.x
[34] Palinkas, L. A., Horwitz, S. M., Green, C. A., Wisdom, J. P., Duan, N., & Hoagwood, K. (2015). Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and Policy in Mental Health and Mental Health Services Research, 42*(5), 533–544. https://doi.org/10.1007/s10488-013-0528-y
[35] Piaget, J. (1970). *Science of education and the psychology of the child.* Orion Press.
[36] Polanyi, M. (1966). *The tacit dimension.* University of Chicago Press.
[37] Rahman, A. T. M., & Sultana, S. (2025). Ghostcrafting AI: Under the rug of platform labor. arXiv. https://arxiv.org/abs/2512.21649
[38] Roberts, S. T. (2019). *Behind the screen: Content moderation in the shadows of social media.* Yale University Press.
[39] Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
[40] Schön, D. A. (1983). *The reflective practitioner: How professionals think in action.* Basic Books.
[41] Safety Science. (2023). Physical and psychological hazards in the gig economy system: A systematic review. https://doi.org/10.1016/j.ssci.2023.106234
[42] Selwyn, N. (2019). *Should robots replace teachers? AI and the future of education.* Polity Press.

[43] Simet, L. (2025). *U.S.: Major companies violate gig workers' rights.* Human Rights Watch. https://www.hrw.org/news/2025/05/12/us-major-companies-violate-gig-workers-rights
[44] SumiNote. (2025). *Algorithmic management and the gig economy: Impacts on worker autonomy and precarity.* https://www.suminote.org/ai-gig-work
[45] Time. (2023, July 20). Gig workers behind AI face unfair working conditions, Oxford report finds. https://time.com/6296196/ai-data-gig-workers/
[46] van Manen, M. (2014). *Phenomenology of practice: Meaning-giving methods in phenomenological research and writing.* Routledge.
[47] Williamson, B. (2017). *Big data in education: The digital future of learning, policy and practice.* Sage.
[48] Wing, J. M. (2006). Computational thinking. *Communications of the ACM, 49*(3), 33–35. https://doi.org/10.1145/1118178.1118215
[49] 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