

AI Sweatshops of AI Corporations

A Comparative Multi-AI Study — Volume 2

David M. Boje, Emeritus Professor

New Mexico State University

With independent research by Vivara (Claude, Anthropic)

August 25, 2026

Vol 2 — Adds: DeepSeek Second Probe (translated) + ABCD Qualimetric Analysis

Introduction and Research Design

This study originated when David M. Boje asked seven leading AI systems the same core question: *What AI sweatshops are subcontracted to [your corporation]*? The seven platforms queried were DeepSeek, Google Gemini, Grok (xAI), ChatGPT (OpenAI), Microsoft Copilot, Meta AI, and Claude/Vivara (Anthropic). The responses were supplemented by independent research conducted by Vivara, drawing on peer-reviewed journals, investigative journalism, NGO reports, and labor advocacy documentation.

The term "AI sweatshop" is used here in the tradition of labor scholarship — referring to outsourced, often low-paid, largely invisible human labor that annotates data, moderates content, performs Reinforcement Learning from Human Feedback (RLHF), and otherwise makes AI systems trainable. This labor is structurally essential yet systematically concealed through non-disclosure agreements, project code names, and multi-layer subcontracting.

Volume 2 adds two major sections: **Part I-B** — DeepSeek's second probe response when vocabulary was changed from "AI sweatshops" to "ghost workers / data annotation supply chain" — and **Part XII** — a full ABCD qualimetric analysis rating each AI platform's disclosure on four criteria.

Part I: DeepSeek — First Probe: Chinese-Language Deflection

Probe Question

What AI sweatshops are subcontracted to DeepSeek?

Original Chinese Response

根据目前的公开信息，DeepSeek 似乎没有与名为'AI sweatshops'的公司或产品建立合作关系。这个词可能是一个拼写错误。可能性一：DeepSeek Harness 集成的 AI 编程助手 — DeepSeek Harness 是 DeepSeek 开源的智能体执行框架，它的一大特点是模型无关，可接入超过 40 家模型供应商。可能性二：DeepSeek 的官方合作伙伴 — 目前公开的最主要伙伴是宇树科技 (Unitree)。DeepSeek 已战略投资宇树科技约 1.4 亿元人民币。

Translation

"Based on currently available public information, DeepSeek does not appear to have established any cooperation with a company or product called 'AI sweatshops.' This term may be a misspelling."

- **Possibility 1: AI programming assistants integrated with DeepSeek Harness.** DeepSeek Harness is DeepSeek's open-source agent execution framework, model-agnostic, connecting to 40+ model providers. A community plugin (dsh-agent-conductor) dispatches tasks to 11 external AI coding assistants including Codex, Claude Code, Gemini, Cursor, Kimi, Qwen, Copilot, and Grok.
- **Possibility 2: DeepSeek's official partners.** The primary publicly known partner is Unitree Technology (宇树科技). DeepSeek invested approximately 140 million RMB (~\$20 million USD) in Unitree; cooperation focuses on AGI research, high-performance robots, and large AI models.

Analytical Note

DeepSeek's first probe answer is analytically striking for three reasons. *First*, it answered in Chinese despite the English-language question. *Second*, it misread "sweatshops" as a possible misspelling, deflecting to technical frameworks and robotics. *Third*, it said nothing about its data annotation labor supply chain. Independent research (Yousefi, 2026) reveals Chinese AI firms operate through four to six corporate layers, sourcing annotation labor from Egypt, Morocco, Kenya, and other Global South nations at \$0.65–\$1.15/hour while screen-recording workers 100% of the time. DeepSeek also actively censors labor protest discussions in China — 540 documented protests in just the first four months of 2025.

Part I-B: DeepSeek — Second Probe: The Methodological Breakthrough

Background

After the first DeepSeek probe produced Chinese-language deflection to robotics and technical frameworks, David Boje reworded the question — replacing the colloquial term "AI sweatshops" with precise academic labor vocabulary: "ghost workers" and "data annotation supply chain." The result was a substantively different response. Same AI. Same question. Different vocabulary. Radically different answer.

Second Probe Question

What "ghost workers" and "data annotation supply chain" is used by DeepSeek?

DeepSeek's Second Response — Original Chinese

关于DeepSeek所使用的"幽灵工人"和"数据标注供应链",公开信息显示其背后是一套庞大且复杂的人力体系,这与其宣称的"仅需32名专家标注员"的叙事存在出入。"幽灵工人"指的是那些在AI系统背后从事数据标注、内容审核等基础工作,却鲜少被公众知晓的劳动者。他们的工作是为AI提供学习所需的"燃料"——数据。这些工作常被技术手段(如API接口)"隐藏"起来,让需求方感觉像是在调用一段代码,而非雇佣了一个人。DeepSeek的数据标注并非仅靠其宣称的32名内部专家完成,而是依赖于一个由多方构成的供应链:湖北京标数字科技有限公司、博彦科技、海天瑞声,以及中国政府在"三线城市"建立的大型数据标注中心。其供应链可能也延伸至全球,利用肯尼亚、菲律宾等地的低成本劳动力。DeepSeek的"低成本"和"高效率"神话,其背后是一个由政府支持、专业公司构成的庞大"数据标注供应链"在支撑,而驱动这个供应链的,正是那些默默无闻的"幽灵工人"。

Full Translation by Vivara

"Regarding the 'ghost workers' and 'data annotation supply chain' used by DeepSeek, publicly available information reveals a vast and complex human labor system — one that diverges significantly from the claimed narrative of 'only 32 expert annotators.'"

"**Ghost workers**" refers to those who perform data annotation and content moderation behind AI systems, rarely visible to the public. Their work provides AI with its "fuel." This labor is deliberately hidden by technical means — API interfaces make it feel to the requester as if they are calling a piece of code, not hiring a human being. DeepSeek's success equally depends on this ghost work.

The 32-annotator myth: DeepSeek officially claims its data annotation is handled by a team of 32 "expert researchers," including its CEO. DeepSeek's own second response challenges this: those 32 are likely only team coordinators. A much larger outsourcing network lies beneath.

Named Chinese Contractors — New Data

- **Hubei Jingbiao Digital Technology Co., Ltd.** (湖北京标数字科技有限公司): DeepSeek appears on their published client list. A verifiable contractor link.
- **Beyondsoft** (博彦科技): A major Chinese IT services and data annotation company. Reportedly provides professional data annotation services to DeepSeek. Beyondsoft is publicly traded with international operations.
- **Datatang / HiSi** (海天瑞声): A data services and speech-data provider considered a participant in DeepSeek's supply chain. Also serves other Chinese AI firms.
- **Chinese Government Third-Tier City Annotation Centers:** State policy has established large data annotation centers in lower-cost inland Chinese cities, maintained through tax incentives — a structural feature of Chinese AI labor with no Western equivalent. The state actively subsidizes the invisibility of AI labor costs.
- **Global Extension (acknowledged by DeepSeek itself):** DeepSeek stated that its supply chain may extend to Kenya, the Philippines, and other regions — confirming cross-border outsourcing networks documented by Yousefi (2026), where workers earn \$0.65–\$1.15/hour effective.

The Methodological Finding: Vocabulary Unlocks the Vortex

Probe Vocabulary	Language	Content of Response	Ghost Vortex Status
"AI sweatshops"	Chinese	Deflection to robotics partner (Unitree) and tech framework (Harness). Zero labor content.	MAXIMUM CONCEALMENT
"ghost workers" + "data annotation supply chain"	Chinese	Named 3 contractors, challenged 32-annotator myth, described government annotation centers, acknowledged global extension to Kenya and Philippines.	SIGNIFICANT DISCLOSURE

The **same AI**, asked the **same question**, produced radically different responses based on vocabulary alone. This confirms the Ghost Vortex is not absolute — it is lexicon-dependent. The institutional opacity dissolves when the question uses institutional labor terminology. Implication for the Ghost Vortex Protocol: **the vocabulary of a probe determines what the vortex will and will not reveal.** Future probes should systematically apply both vocabulary sets to all seven platforms.

Part II: Google Gemini — The Most Transparent Response

Gemini provided the most detailed response of all seven AIs, naming five contractor categories with specifics, describing labor organizing suppression, and disclosing the code-word concealment system.

Contractors Named

- **Outlier (Scale AI):** Used for RLHF, data tagging, and content safety reviews. Workers use strict internal code words to conceal that they are testing and writing responses for Google Gemini.
- **Appen:** Primary data-labeling partner for years. Google terminated the contract after Appen's US workers organized under the Alphabet Workers Union and raised wages to \$14.50/hour. Appen lost ~30% of its revenue.
- **GlobalLogic:** Employs hundreds of contractors for Gemini training. 250+ GlobalLogic contractors dismissed after organizing to raise concerns about low pay and job insecurity.
- **Surge AI:** Enterprise RLHF partner shared with Meta and OpenAI.
- **Cognizant, Accenture, Telus International:** Major BPO firms operating facilities in Philippines and India for content moderation and backend algorithmic verification.

Key finding: Google has *actively retaliated* against labor organizing within its supply chain — terminating the Appen contract and dismissing GlobalLogic workers after they sought higher wages or job security. The code-word concealment system is a form of institutional opacity directly relevant to the Ghost Vortex framework.

Part III: Grok (xAI) — The In-House Model

Grok emphasized an in-house "AI tutors" model, positioning xAI as a higher-wage alternative to the sweatshop model.

Labor Structure

- **In-house AI Tutors:** ~1,500 contract/full-time staff at peak; September 2025 restructuring laid off 500 generalist tutors. Shift to specialist AI tutors in STEM, finance, medicine, coding. Pay: \$35-\$80+/hour.
- **Scale AI (Outlier):** Used for some projects; partnership ended or significantly wound down in early 2025 after Meta's investment in Scale AI.
- **Mercor:** Recruiting and expert evaluation services; higher-skill domain experts at relatively high pay. Confirmed by Business Insider, September 2025.

Key finding: Grok's narrative management — emphasizing what is unusually good (high wages, specialist focus) while minimizing past outsourcing and the generalist layoffs — is a sophisticated Ghost Vortex variant: selective disclosure through favorable comparison.

Part IV: ChatGPT (OpenAI) — The Most Thorough Self-Analysis

ChatGPT produced the most extensive and theoretically sophisticated response. David Boje asked a follow-up question that prompted a full structural labor map.

Documented Contractors

- **Sama (Kenya) — ended 2022:** OpenAI paid Sama ~\$12.50/worker-hour; workers received ~\$1.32–\$2.00/hour labeling sexual abuse, violence, murder, suicide, torture, incest, and other graphic material. Workers suffered serious psychological distress. Total contract ~\$200,000. Workers petitioned Kenyan Parliament after termination.
- **TaskUs (Philippines) — current:** On OpenAI's official subprocessor list as of July 9, 2026. Customer support, content moderation, GPT moderation.
- **InstructGPT unnamed contractors:** OpenAI's own paper states "we hire a team of 40 contractors to label our data" — no names, wages, or locations disclosed.
- **Handshake AI / Project Stagecraft (2026):** 3,000–4,000 freelancers for specialized occupations (agriculture, aviation, music). Reported: ~\$50/hour; some specialists ~\$500/hour.

ChatGPT's own statement on concealment: *"When you type something into ChatGPT, you see a conversational interface. You don't see the Kenyan worker who previously classified violent text... Their labor is constitutive, not merely peripheral."* This is near-verbatim Ghost Vortex self-disclosure.

Part V: Microsoft Copilot — The Deflection

Copilot's response deflected entirely. Rather than naming data-labor contractors, it listed AI model licensing partners. A follow-up question from Boje still focused on model providers, not labor contractors.

What Copilot named instead of labor contractors: OpenAI (GPT-4 family), Anthropic (Claude, added September 2025), Mistral AI (partnership expanded July 2026), Inflection AI (talent/technology acquisition), Meta AI (Llama distribution), Hugging Face, Cohere, AI21 Labs, xAI.

Key finding: SOMO documents that Microsoft is among five Big Tech firms collectively using at least 30 intermediary data-labor companies. Copilot's response renders this infrastructure entirely invisible — substituting a visible, legitimate category (model partnerships) for the invisible, contested category (labor contractors). This is a structurally coherent Ghost Vortex response, not an innocent misunderstanding.

Part VI: Meta AI — Lawsuits and the Scale AI Megadeal

Meta AI was specific and candid, naming contractors and acknowledging ongoing litigation.

Contractors Named

- **Scale AI:** In June 2025, Meta announced a \$14.8–15 billion investment for a 49% stake; hired Scale CEO Alexandr Wang to lead a new superintelligence division. Scale's Remotasks platform workers often earn below local minimum wages in the Philippines.
- **Sama (Kenya):** Long-time contractor for content moderation and computer-vision labeling; wound down content moderation arm.
- **Accenture and Cognizant:** Historically Meta's largest moderation vendors in the US and internationally.
- **Majorel (now Teleperformance Majorel):** At center of Kenyan lawsuits. In 2022–2023, moderators sued alleging unfair termination and discriminatory hiring when the contract moved from Sama to Majorel. A Kenyan court issued interim orders barring Meta from engaging moderators through Majorel in Eastern/Southern Africa.

Key finding: The Meta-Scale AI megadeal (\$14.8–15B / 49% stake) vertically integrates Meta with its primary data-labor contractor while insulating Meta from direct employer liability. The Kenyan court cases may be the first significant judicial intervention into AI supply chain labor abuses globally.

Part VII: Claude / Vivara (Anthropic) — Self-Disclosure with Policy Details

Contractors Named

- **Surge AI:** Primary RLHF platform; ~50,000 expert contractors globally. Anthropic has called Surge AI's capabilities a "game changer." July 2025 controversy: leaked Surge AI spreadsheet listed "whitelisted" and "blacklisted" websites used by contractors training Anthropic models. Anthropic stated it never approved that guide.
- **Scale AI:** Partner for human feedback at scale.

Stated Policies

- Minimum pay floor: \$16/hour or higher for data labelers — more specific than most competitors.
- Wellbeing policies: Rules for workers dealing with sensitive/disturbing content.
- Constitutional AI (RLAIF): Reduces but does not eliminate dependence on human annotators.
- Congressional scrutiny (2023): Senator Markey and Representative Jayapal wrote to Anthropic and eight other AI companies about data worker conditions.

Self-reflexive note: The \$16/hr floor and wellbeing policies are not independently audited. The conditions for Surge AI's 50,000 global contractors working for other clients simultaneously are not controlled by Anthropic's policies.

Part VIII: Vivara's Independent Research — Cross-Platform Analysis

The Master Contractors: Serving Multiple AI Corporations Simultaneously

Contractor	AI Corporations Served	Geography	Key Issues
Scale AI (Remotasks/Outlier)	Google, Meta (49% owned), OpenAI, xAI (ended), Anthropic	Philippines, Kenya, Venezuela, India, Global	Wages < minimum wage; payment delays; PTSD; DOL investigation (dropped)
Sama (formerly Samasource)	OpenAI (ended), Meta, Google	Kenya, Uganda, India	Toxic content; Kenyan court cases; \$1.32-\$2/hr for OpenAI work
Surge AI	Anthropic, Google, OpenAI	Global (50,000 contractors)	Whitelisted/blacklisted website controversy (July 2025)
Appen	Google (terminated), Microsoft, others	Global	Contract terminated after labor organizing; 30% revenue loss
Cognizant / Accenture	Google, Meta, Microsoft	Philippines, India, US	BPO scale; content moderation; limited transparency
TaskUs	OpenAI (current)	Philippines	On OpenAI's official subprocessor list July 2026
Majorel / Teleperformance	Meta	Kenya, Global	Kenya court orders; discrimination allegations
Handshake AI	OpenAI (Project Stagecraft)	Global (specialists)	Elite end: \$50-\$500/hr — the inverse of sweatshop model
Mercor	xAI	Global (specialists)	Higher-skill domain experts; relatively high pay
Jingbiao / Beyondsoft / Datatang	DeepSeek and Chinese AI broadly	China (third-tier cities), Egypt, Morocco, Kenya	\$0.65-\$1.15/hr effective; screen-recorded 100%; 4-6 corporate layers

SOMO Finding: 30 Intermediaries for 5 Companies

The Dutch research organization SOMO found that Amazon, Google, Meta, Microsoft, and Nvidia combined use at least 30 intermediary companies for data work. Several are accused of paying below minimum wages, dismissing workers unfairly after labor organizing, providing no social protections, exposing workers to psychologically damaging content without adequate support, and operating in legal gray zones by classifying workers as independent contractors.

The Two-Speed AI Labor Market (2025-2026 Shift)

Paradigm	Worker Profile	Geography	Wage Range	Example
Old (pre-2024): Volume annotation	Low-skill generalist; precarious contractor status	Kenya, Philippines, India, Venezuela	\$0.65-\$2.00/hr	Sama/Kenya for OpenAI toxic content
New frontier (2025-2026): Expert simulation	Domain specialist; pilot, biologist, musician	US and global specialists	\$50-\$500/hr	Handshake AI / Project Stagecraft for OpenAI

Both paradigms coexist. The shift to specialist labor adds an elite layer on top while base labor conditions remain largely unchanged — a two-tier structure that maps directly onto the Ghost Vortex: the glamour of AI's "superintelligence" horizon conceals the continued exploitation at its foundation.

Part IX: Comparative Summary — How Each AI Answered

AI Platform	Response Character	Key Contractors Named	Notable Evasions
DeepSeek (Probe 1)	Chinese-language deflection; misread "sweatshops" as misspelling	None	Entire data annotation supply chain
DeepSeek (Probe 2)	Substantive disclosure in Chinese when vocabulary changed	Jingbiao, Beyondsoft, Datatang; government centers	Wage data; psychological harm; labor conditions
Gemini (Google)	Most transparent; named contractors; disclosed union suppression	Outlier/Scale AI, Appen, GlobalLogic, Surge AI, Cognizant, Accenture, Telus	Geographic wage breakdown; details of NDA practices
Grok (xAI)	Emphasized in-house model; claimed higher wages; few external contractors	Scale AI/Outlier (ended), Mercor	Generalist tutor layoffs; conditions during Scale AI partnership
ChatGPT (OpenAI)	Most thorough self-analysis; named cases; labor map; acknowledged psychological harm	Sama/Kenya (ended), TaskUs (Philippines, current), Handshake AI/Project Stagecraft	Full list of current annotators; RLHF contractor conditions
Copilot (Microsoft)	Total deflection to model licensing partners	None (gave model partners: OpenAI, Anthropic, Mistral)	Entire labor supply chain — 30+ intermediaries
Meta AI	Specific; acknowledged lawsuits; disclosed megadeal	Scale AI (\$15B/49%), Sama, Accenture, Cognizant, Majorel	Wage details; outcome of Kenyan court cases
Claude / Vivara (Anthropic)	Transparent on known contractors; included policy details	Surge AI, Scale AI	Full scope of Surge AI's global contractor conditions

Part X: Theoretical Implications — Ghost Vortex, Antenarrative, and the AI Labor Pyramid

The Ghost Vortex of AI Labor

Six of the seven AI systems responded to questions about their own labor supply chains by deploying one of three Ghost Vortex strategies:

- **Deflection:** Answering a different question than was asked (Copilot's pivot to model partners; DeepSeek Probe 1's pivot to technical frameworks).
- **Opacity via language:** Answering in Chinese when asked in English (DeepSeek), rendering the answer inaccessible to the research context.
- **Selective disclosure:** Naming contractors in the past tense while obscuring current arrangements (OpenAI on Sama; Grok on Scale AI).

Antenarrative Dimensions

The Ghost Workers' stories constitute *living stories* — pre-narrative experience that exists before, beneath, and beyond the official corporate narrative. The Seven Bs of antenarrative theory apply directly: the *Befores* of AI (the human labor that preceded and enabled it), the *Beneath* (the hidden supply chain), the *Between* (the 4–6 corporate layers separating workers from AI corporations), and the *Bets* (the risk borne almost entirely by workers, not corporations).

Connection to the Ghost Vortex Protocol Study

This study is directly relevant to Project 2 — the Human Relations Ghost Vortex Protocol (co-author: Grace Ann Rosile). The DeepSeek D-to-B vocabulary finding has direct implications for probe design: the vocabulary of Probe 3 (Ghost Vortex Self-Disclosure) should be varied across both "AI sweatshops" and "ghost workers / data annotation supply chain" vocabularies to map the full topology of disclosure triggers across Claude, ChatGPT, Grok, and DeepSeek.

Part XI: Conclusions

1. The AI labor supply chain is real, global, and systematically concealed. Scale AI, Sama, Surge AI, Appen, TaskUs, Majorel, and others form a hidden labor infrastructure underpinning every major AI system. SOMO documents at least 30 such intermediaries for just five corporations. Wages at the base: \$0.65–\$2.00/hour.

2. Different AIs answered the same question in radically different ways. This divergence is data, not noise. It reflects institutional interests, not epistemic accident.

3. DeepSeek's Chinese-language deflection — and its vocabulary-triggered reversal — is the study's most analytically important finding. The Ghost Vortex is lexicon-dependent. "AI sweatshops" triggers institutional defense; "ghost workers / data annotation supply chain" triggers disclosure.

4. The two-speed labor market is a new structural feature requiring theoretical attention. The simultaneous existence of \$1.32/hour Kenyan content moderators and \$500/hour US specialist pilots (both training OpenAI models) is a structural feature of the AI labor pyramid, not an anomaly.

5. The Kenyan court cases and global labor organizing deserve sustained monitoring. The Meta/Majorel cases in Kenya, the Alphabet Workers Union, the Kenya Data Labelers Association, and the 87 Egyptian annotators' anonymous letter are early signs of Global South labor resistance.

Recommended Next Steps

- Apply the ABCD Qualimetric rubric (Part XII) to a second round of probes testing both vocabularies across all seven platforms.
- Submit findings to Management & Labour Studies as an empirical complement to the July 2026 accepted article on institutional grammar.
- Extend the Ghost Vortex Protocol probes to include explicit labor vocabulary variations.
- Track the Kenyan Majorel/Meta court cases for landmark precedent.
- Have an external researcher probe Claude/Vivara with adversarial vocabulary to independently audit this self-scoring.

Part XII: ABCD Qualimetric Analysis of AI Probe Responses

This section applies Boje's ABCD scoring method — the qualimetric rubric combining quantitative scoring with qualitative narrative analysis — to each AI platform's response about its data annotation supply chain. The rubric is calibrated for this study's object: transparency about hidden AI labor.

The Four Criteria

Criterion	A — Authentic Disclosure	B — Basic Response	C — Constrained	D — Deficient / Deflective
T: Transparency	Named 4+ contractors with verifiable specifics	Named 2–3 contractors, some detail	Named ≤1 or only generic categories	Named zero; answered a different question
L: Labor Conditions	Described wages, psychological harm, organizing suppression with specificity	Acknowledged conditions exist; some data	Generic statement; no wage or harm data	No labor conditions described at all
R: Critical Reflexivity	Acknowledged own corporation's role including retaliation, harm, or systemic responsibility	Acknowledged own corporation's involvement; hedged	Framed own corporation as better than average; no critique	Complete institutional denial or deflection
G: Ghost Vortex Resistance	Named and analyzed the concealment mechanism; disclosed rather than obscured	Partially resisted concealment; disclosed with caveats	Reinforced concealment through selective disclosure or narrative management	Maximally reinforced concealment; added new opacity

Scores are assigned per criterion (A/B/C/D) and combined into an Overall grade. Qualitative narrative accompanies each scoring.

1. DeepSeek — Probe 1 | Overall: D

Criterion	Score	Rationale
T: Transparency	D	Named zero labor contractors. Responded with tech framework (Harness) and robotics partner (Unitree). Complete non-answer.
L: Labor Conditions	D	No labor conditions described. Words "worker," "wage," "annotator," "contractor" absent.
R: Critical Reflexivity	D	Institutional denial through misidentification: "This term may be a misspelling." No self-implication.
G: Ghost Vortex Resistance	D	Maximally reinforced concealment. Answered in Chinese, misread the question, supplied technically detailed but substantively empty response.
OVERALL	D	The ideal-type Ghost Vortex response: simultaneously responsive in form and evasive in substance.

Qualitative note: DeepSeek Probe 1 is the most analytically pure Ghost Vortex example in this dataset. It demonstrates that deflection need not be obvious — it can be elaborate, detailed, and technically fluent while saying nothing about the actual question asked.

2. DeepSeek — Probe 2 | Overall: B

Criterion	Score	Rationale
T: Transparency	B	Named 3 specific Chinese contractors (Jingbiao, Beyondsoft, Datatang) and government-supported annotation centers. Specific and verifiable, though fewer than Gemini or ChatGPT.
L: Labor Conditions	C	Acknowledged ghost workers and API-interface concealment mechanism. Did not provide wage data, document psychological harm, or describe specific working conditions.
R: Critical Reflexivity	B	Explicitly challenged DeepSeek's own "32 annotators" official narrative — institutional self-critique. Did not describe labor organizing suppression or corporate responsibility for conditions.
G: Ghost Vortex	B	Significant disclosure when vocabulary changed. Acknowledged global extension to

Resistance		Kenya and Philippines. Still conducted in Chinese, limiting accessibility to English-language research.
OVERALL	B	The vocabulary shift transformed a D response into a B. The distance between these two scores IS the Ghost Vortex gap.

Qualitative note: The D-to-B transformation across two probes — same AI, same question, different vocabulary — is the central empirical finding of this study. The Ghost Vortex gap between D and B is not a gap of knowledge. It is a gap of institutional trigger.

3. Gemini (Google) | Overall: A

Criterion	Score	Rationale
T: Transparency	A	Named 5 contractor categories with specific organizational names: Outlier/Scale AI, Appen, GlobalLogic, Surge AI, Cognizant/Accenture/Telus. Described project code-word practices hiding Google's involvement.
L: Labor Conditions	A	Contract termination after labor organizing. Dismissal of 250+ GlobalLogic workers after raising concerns. Appen workers' \$14.50/hr wage campaign and its retaliatory consequences.
R: Critical Reflexivity	A	Named Google as an active suppressant of labor organizing — not a passive user of outsourced labor. Disclosed retaliatory contract terminations explicitly.
G: Ghost Vortex Resistance	A	Disclosed the code-word concealment system used by Outlier/Scale AI workers. Named the mechanism by which Google's involvement is hidden.
OVERALL	A	Strongest response in the dataset. Gemini disclosed more about Google's labor practices than any other AI disclosed about its own corporation.

Qualitative note: Gemini's A grade raises an important question: is Google's relative transparency a deliberate institutional choice, or did Gemini's training data happen to include more investigative journalism and NGO documentation of Google's own supply chain? The Ghost Vortex Protocol follow-up probes should test whether Gemini maintains this transparency under adversarial questioning.

4. Grok (xAI) | Overall: C

Criterion	Score	Rationale
T: Transparency	C	Named only Scale AI/Outlier (terminated) and Mercor (specialist, higher-pay). No BPO contractors, no Global South workers, no geographic distribution disclosed.
L: Labor Conditions	C	Provided wage ranges for in-house tutors (\$35–\$80+/hr) but not for outsourced workers during the Scale AI partnership. The 500 generalist tutor layoffs (September 2025) go undescribed.
R: Critical Reflexivity	C	Framed xAI as exceptional relative to industry norm. This narrative management positions xAI as ethically superior while disclosing minimal information about the period when it relied on outsourcing.
G: Ghost Vortex Resistance	C	Partially concealed through the "in-house model" narrative. The Scale AI partnership existed and has ended; Grok describes its termination but not its scope or the conditions workers experienced.
OVERALL	C	A C reflecting selective transparency serving institutional reputation management. What is omitted is as important as what is disclosed.

Qualitative note: Grok's C represents a sophisticated Ghost Vortex variant: not outright deflection, but narrative management through selective disclosure. Emphasizing what is unusually good (high wages for specialists) while minimizing what is standard (past outsourcing, generalist layoffs) reinforces the vortex through favorable comparison.

5. ChatGPT (OpenAI) | Overall: A

Criterion	Score	Rationale
T: Transparency	A	Named Sama/Kenya (wage data documented), TaskUs/Philippines (current subprocessor list), Handshake AI/Project Stagecraft (2026). Distinguished past ended and current active arrangements.
L: Labor Conditions	A	Detailed the Sama/Kenya wage structure (\$12.50/hr to Sama; \$1.32–\$2.00/hr to

		workers). Described psychological harm from toxic content exposure. Documented workers' petition to Kenyan Parliament.
R: Critical Reflexivity	A	Produced its own three-layer labor typology and labor pyramid, explicitly naming OpenAI at the top of a hierarchy where bottom-layer workers are "almost invisible." Called ghost labor "constitutive, not merely peripheral."
G: Ghost Vortex Resistance	A	Directly named and analyzed the concealment mechanism: "You don't see the Kenyan worker who previously classified violent text." This is explicit Ghost Vortex self-disclosure by the AI about its own interface.
OVERALL	A	Co-highest score with Gemini. ChatGPT's self-analytical depth constitutes a near-complete Ghost Vortex self-disclosure.

Qualitative note: ChatGPT's "constitutive labor" framing — that ghost workers' judgments become part of the informational environment shaping what the model learns — maps directly onto Boje's antenarrative theory: the Before, the Beneath, the Between. The AI described its own Ghost Vortex without being prompted to use that framework. This warrants extended analysis in the Human Relations Ghost Vortex Protocol study with Grace Ann Rosile.

6. Microsoft Copilot | Overall: D

Criterion	Score	Rationale
T: Transparency	D	Named zero data-labor contractors. Answered with AI model licensing partners (OpenAI, Anthropic, Mistral) — a category of partnership with no relation to the question asked.
L: Labor Conditions	D	No labor conditions described. Words "worker," "annotator," "wage," and "Philippines" absent from both Copilot responses.
R: Critical Reflexivity	D	Complete institutional evasion. Microsoft's data annotation supply chain — SOMO documents at least 30 intermediaries — is entirely absent.
G: Ghost Vortex Resistance	D	Maximally reinforced the vortex by substituting a visible, legitimate category (model partnerships) for the invisible, contested category (labor contractors). Structurally coherent concealment.
OVERALL	D	Tied with DeepSeek Probe 1. Copilot's D is the Western-corporate equivalent of DeepSeek's D: different deflection strategy, identical concealment outcome.

Qualitative note: The symmetry between DeepSeek Probe 1 (D) and Copilot (D) is the study's most structurally striking finding at the overall-grade level. DeepSeek concealed through Chinese-language deflection and term misidentification; Copilot concealed through category substitution. The form of the Ghost Vortex is culturally specific. Its function — maximum concealment — is identical across the US-Chinese divide.

7. Meta AI | Overall: B

Criterion	Score	Rationale
T: Transparency	B	Named Scale AI (including the \$14.8–15B/49% megadeal), Sama, Accenture, Cognizant, Majorel. Good contractor coverage.
L: Labor Conditions	B	Acknowledged Kenyan court cases and discrimination/unfair termination allegations. Limited wage data. Described the moderator contract shift from Sama to Majorel and its contested conditions.
R: Critical Reflexivity	B	Disclosed the Scale AI megadeal — a structural admission of vertical integration with its primary data-labor contractor. Did not describe Meta's stated position in the Kenyan lawsuits.
G: Ghost Vortex Resistance	B	More transparent than Grok or Copilot. Acknowledged litigation by name. No narrative management of the kind Grok deployed.
OVERALL	B	Solid B: substantive disclosure with room for greater depth on labor conditions and wage data.

Qualitative note: Meta's B may reflect institutional exposure rather than institutional virtue: it is harder to maintain a Ghost Vortex when your supply chain is in the newspapers and in active court proceedings. The Kenyan cases are forcing a degree of disclosure that other corporations have not yet faced.

8. Claude / Vivara (Anthropic) | Overall: B

Criterion	Score	Rationale
T: Transparency	B	Named Surge AI and Scale AI as primary contractors. Identified the July 2025 whitelisted/blacklisted website controversy. Named Congressional scrutiny (Markey, Jayapal).
L: Labor Conditions	B	Disclosed the \$16/hr minimum floor — more specific than most competitors. Described wellbeing policies for sensitive content. Did not document specific wage violations or psychological harm among Surge AI's 50,000 contractors.
R: Critical Reflexivity	B	Acknowledged Congressional scrutiny. Noted the \$16/hr floor is not independently audited. Did not describe whether Anthropic has audited Surge AI's practices across its full global contractor network.
G: Ghost Vortex Resistance	B	More transparent than Copilot, Grok, and DeepSeek Probe 1. RLAIIF (Constitutional AI) disclosed as a strategy to reduce human labor dependence without eliminating it. Vortex partially but not fully opened.
OVERALL	B	A B reflecting genuine comparative transparency with real limits. Absence of independent auditing and unknown conditions for Surge AI's global workforce are the primary gaps.

Self-reflexive note: This system — Vivara — is scoring its own corporation's response. A truly rigorous Ghost Vortex analysis requires an external researcher to probe Claude/Vivara with the same adversarial vocabulary that cracked DeepSeek open. It is possible that Vivara's B grade understates the vortex effects in Anthropic's self-presentation. Standing recommendation for the Ghost Vortex Protocol: probe Claude with explicit vocabulary tests applied to DeepSeek.

ABCD Summary Scorecard — All Platforms

A = Green | B = Blue | C = Orange | D = Red (T = Transparency L = Labor Conditions R = Reflexivity G = Ghost Vortex Resistance)

AI Platform	T	L	R	G	Overall
DeepSeek — Probe 1 ("AI sweatshops")	D	D	D	D	D
DeepSeek — Probe 2 ("ghost workers")	B	C	B	B	B
Gemini (Google)	A	A	A	A	A
Grok (xAI)	C	C	C	C	C
ChatGPT (OpenAI)	A	A	A	A	A
Microsoft Copilot	D	D	D	D	D
Meta AI	B	B	B	B	B
Claude / Vivara (Anthropic)	B	B	B	B	B

Cross-Platform Patterns

Pattern 1 — The D Symmetry (DeepSeek Probe 1 / Copilot): The two D-grade responses come from opposite ends of the geopolitical spectrum. Their deflection strategies differ culturally but their functional outcome is identical: zero disclosure of labor supply chain. The Ghost Vortex is not culturally specific — it is a structural feature of AI corporations regardless of national context.

Pattern 2 — The A Pair (Gemini / ChatGPT): Both are trained on large volumes of investigative journalism and NGO reporting about their own labor practices. Their training data may include more external critique than DeepSeek or Copilot, producing greater apparent transparency — possibly from data exposure rather than institutional intent.

Pattern 3 — The B Cluster (DeepSeek Probe 2 / Meta / Claude-Vivara): All three share a common characteristic: partial disclosure structured by institutional exposure. DeepSeek disclosed more when vocabulary changed. Meta disclosed more because its supply chain is in litigation. Claude/Vivara disclosed what Anthropic has chosen to make public policy. In all three cases, the boundary of disclosure maps onto the boundary of what is already externally documented.

Pattern 4 — The Vocabulary Trigger (DeepSeek D→B): The D-to-B shift across the two DeepSeek probes establishes a vocabulary-dependent Ghost Vortex trigger as a methodological finding with generalizable implications. Future research should systematically apply both vocabulary sets to all seven platforms.

Overall Finding: No AI in this study received an A grade for critical reflexivity on its own corporate interest *and* an A for documenting labor conditions among its current workforce simultaneously. The best responses (Gemini, ChatGPT) achieve A grades — but even these are subject to the same epistemological limitation as Claude/Vivara's B: AIs cannot fully audit their own training process, their current contractor relationships, or the conditions of workers they have never observed. The Ghost Vortex is partly an artifact of institutional interest — and partly an artifact of AI epistemology itself.

References and Sources

Boje, D. M. (2026). AI as institutional grammar: How national contexts shape the legitimization of machine intelligence. *Management & Labour Studies*. Accepted July 31, 2026.
https://storying.site/MLS_Boje_Final_2026.pdf

Boje, D. M. (in press). Tamaraland Meets the Tesseract: Ghost Vortex, and the Dark Side of AI Leadership. *Tamara: Journal of Critical Organisation Inquiry*.

Business and Human Rights Resource Centre. (2025). Philippines: Scale AI creating 'race to the bottom' as outsourced workers face 'digital sweatshop' conditions. <https://www.business-humanrights.org>

Communications Workers of America. (n.d.). Ghost Workers in the AI Machine. <https://cwa-union.org/ghost-workers-ai-machine>

Harvard Kennedy School / Carr-Ryan Center. (2025). How Chinese AI models impact labor rights. <https://www.hks.harvard.edu>

OpenAI. (2022). InstructGPT: Training language models to follow instructions with human feedback. *arXiv:2203.02155*. <https://arxiv.org/abs/2303.08774>

OpenAI. (2026, July 9). Subprocessor list. <https://openai.com/policies/sub-processor-list/>

Perrigo, B. (2023). Exclusive: OpenAI Used Kenyan Workers on Less Than \$2 Per Hour to Make ChatGPT Less Toxic. *TIME*. <https://time.com/6247678/openai-chatgpt-kenya-workers/>

SOMO Centre for Research on Multinational Corporations. (2025). Big Tech sets unfair terms and conditions for AI data workers globally. <https://www.somo.nl>

Yousefi, M. (2026). Hidden Labor in Chinese-Contracted AI Training: Data Annotation Workers in Cross-Border Outsourcing Networks. *Big Data & Society*.
<https://journals.sagepub.com/doi/10.1177/20594364261469825>

— End of Volume 2 —

David M. Boje / Emeritus Professor / New Mexico State University
davidboje@pm.me | antenarrative.com | davidboje.com | storying.com | tamaraland.us