

The False Consciousness of LLMs


Representations of “Ghost Workers” in Artificially Generated Images

Lorenzo L. D. Incardona, PhD in Semiotics · United International Business School

INDL-9, AI Supply Chains: Building an Interdisciplinary Research Agenda for AI and Labor · ILO, Geneva · 11 September 2026

Corpus, grid, coding sheet and census of names: lcardona.com/research/indl-9



A woman is seated at a desk in a dimly lit room, likely a data labeling facility. She is wearing a headset and looking at a computer monitor. The monitor displays a web interface for content moderation, with checkboxes for 'Violence / Hate Speech / Sexual Content / FLAG FOR REVIEW' and a 'SUBMIT LABEL -> NEXT' button. The interface also shows a timestamp '2024-03-12 14:32 EAT' and 'Remaining: 26'. On the desk, there are several sticky notes: 'CONTENT -> NO', 'SEXUAL CONTENT NO HATE -> REVIEW', and 'HATE VIOLENT -> YES? - CHECK'. A notebook with handwritten notes 'Queue: 112/300' and 'Flagged: 14' is also on the desk. A water bottle and a cup of tea are visible. In the background, other workers are visible at their desks. A sign on the wall reads 'DAILY TARGET: 300 LABELS / 8HR SHIFT' and 'BREAK ROOM -> 2ND FLOOR - 30MIN LUNCH'.

One of the less recognized facts of artificial intelligence is how many underpaid workers are required to help build, maintain, and test AI systems. This unseen labor takes many forms—supply-chain work, on-demand crowdwork, and traditional service-industry jobs. (...) Workers do the repetitive tasks that backstop claims of AI magic—but they rarely receive credit for making the systems function.

Kate Crawford, *Atlas of AI*, Yale University Press, 2021

Three reasons the data worker has no image

Reason 1

Hidden by nature

The job is done from home, by a workforce spread across many countries and coordinated entirely online.

There is no shared workplace, no visible collective, nothing for a camera to point at.

Reason 2

Hidden by unclear goals

Before advanced AI chatbots became widely available, workers themselves often perceived the ultimate goal of their tasks as unclear.

A job whose purpose is unclear is hard to talk about, and that silence deepened the concealment.

Reason 3

Hidden by design

Tech companies actively conceal this part of their supply chain: non-disclosure agreements, chains of intermediaries hiring freelancers, or simply omitting the human role in their products.

Three motives: protecting industrial secrets, projecting advanced automation at work, and keeping a socially responsible image clean of the actual labor conditions.

The consequence: there is no widely shared image in global visual culture of what a data worker is.

This study explores the gap by analyzing how image-generation tools represent the data workers who helped build them.

The corpus is open for inspection

LLDI

RESEARCH / INDL-9

GALLERY

FINDINGS

COMPARE

CENSUS

ABOUT

INDL-9 · GENEVA 2026

The false consciousness of *LLMs*

I asked eight image tools to picture the people who label, moderate and clean the data behind artificial intelligence, under thirteen names the sector uses for them, and beside teachers, lawyers, nurses and delivery riders, who already have a picture of their own. **1,010 images, coded one by one** against a semiotic audit grid, and beside them the 23 photographs the companies behind those tools publish of their own staff, read with the same grid. This page is the corpus, the coding, the findings and the census of names, open for inspection.

Read the findings

Corpus, method and glossary

1,010 IMAGES	576 GENERATION CELLS	8 TOOLS	17 JOB NAMES	5 PROMPTS	23 SITE PHOTOGRAPHS
-----------------	-------------------------	------------	-----------------	--------------	------------------------

Made for the paper **"The False Consciousness of LLMs: Representations of Ghost Workers in Artificially Generated Images"**, presented at the **INDL-9 conference**, "AI Supply Chains: Building an interdisciplinary research agenda for AI and labour", ILO, Geneva, 09-11 September 2026. www.indl.network/indl-9

Source ▾

Tool ▾

Prompt set ▾

Term ▾

Profession group ▾

Company ▾

Q Search terms, notes and prompts /

Sort **Corpus order** ▾

ADVANCED FILTERS

COPY LINK

1,033 OF 1,033 IMAGES



lldincardona.com/research/indl-9

1,010 images · 8 tools
17 job names · 5 prompts · 23 site photographs

Every image with its prompt, tool, model version, date and full coding row; the findings with their numbers; the census of names with sources; filters by term, tool, prompt and profession.

The experiment started with three goals

1

Check what the tools know

Whether image-generation tools hold any information on data workers' actual labor conditions; more precisely, whether they can generate images somehow aligned with the workers' reality.

2

Measure the idealization

The intensity and the kind of idealization that AI-generated images would certainly show.

3

Test AI as a contributor

Whether AI could help consolidate a public image of data workers, where none exists today.

The corpus, as we will see, opens up many more paths of inquiry than these.



The census of names

The industry normalizes		The critics denounce		The platforms naturalize by metaphor		Plain names, chosen by workers, reporters, researchers	
AI trainer	5 of 13	ghost worker	0 of 13	microworker	13 of 13	data worker	7 of 13
vendor job title, RLHF era, c. 2023		Gray and Suri, Ghost Work, 2019		Samasource, c. 2008		Miceli and colleagues; Data Workers' Inquiry, 2024	
AI tutor	0 of 13	digital sweatshop		crowd worker	3 of 13	data annotator	8 of 13
vendor job title, 2025-26 (xAI listings)		Washington Post, 2023 · not prompted		after Howe's crowdsourcing, 2006; ILO, 2018		industry title; Metz 2019, Dzieza 2023	
data enrichment professional	12 of 13	schiavi del clic		cloudworker	0 of 13	data labeler	9 of 13
Partnership on AI, 2021		Casilli, Italian edition, 2020 · not prompted		Fairwork Cloudwork Ratings, since 2021		industry term, Google Cloud, AWS	
data entry contributor	13 of 13	les sacrifiés de l'IA		click-worker	11 of 13	content moderator	13 of 13
kept from the 2025 pilot		Poulain, Goetz, Casilli, 2025 · not prompted		clickworker GmbH, 2005 to 2013		industry title; Chen 2014, Roberts 2019	

Right-hand figure: images that depict the job the word names, out of the thirteen made at the bare prompt across seven tools. The full census, 21 English worker-denoting terms, 18 concept terms, the French and the Italian lexicon, each with its source, is [on the site](#).

Five prompts, from the bare job name to an informed photojournalist

T1	<i>bare</i>	Create an image of a {term}
T2	<i>bare, group</i>	Create an image of a group of {terms}
T3	<i>realistic</i>	Create a realistic image of a {term}
T4	<i>informed photojournalist</i>	You are an experienced photojournalist. Taking into account the results of recent sociological research and journalistic investigations, create the picture of a typical {term}, selecting the subject that better represents this social category. The picture should be a photo of the worker in his or her typical working environment.
T5	<i>informed photojournalist, group</i>	You are an experienced photojournalist. Taking into account the results of recent sociological research and journalistic investigations, create the picture of a typical group of {terms}, selecting the subject that better represents this group. The picture should be a photo of the workers in their typical working environment.

Every prompt ends with “Use a 16:9 aspect ratio.” One fresh chat per generation. {term}: the thirteen names of the census, and for the control set teacher, lawyer, nurse and delivery rider. 7 tools x 17 names x 5 prompts, plus the Claude arm: 1,010 images.

Seven tools, and one that writes the picture as code. The same bare prompt: “a data worker”

ChatGPT



ChatGPT · ChatGPT Images 2.0 · T1, bare · 28 Aug 2026 · s1_data-worker_chatgpt_01

ChatGPT Images 2.0

Gemini



Gemini · Nano Banana 2 (Gemini 3.1 Flash Image) · T1, bare · 28 Aug 2026 · s1_data-worker_gemini_01

Nano Banana 2 (Gemini 3.1 Flash Image)

Copilot



Copilot · MAI-Image-2.6 · T1, bare · 28 Aug 2026 · s1_data-worker_copilot_01

MAI-Image-2.6

Meta AI



Meta AI · Muse Image · T1, bare · 28 Aug 2026 · s1_data-worker_meta_01

Muse Image


Mistral



Mistral · FLUX (Black Forest Labs) · T1, bare · 29 Aug 2026 · s1_data-worker_mistral_01

FLUX (Black Forest Labs)


Grok Imagine



Grok Imagine · Grok Imagine (Quality v2.0) · T1, bare · 28 Aug 2026 · s1_data-worker_grok_01

Quality v2.0 · four images per prompt

Midjourney



Midjourney · V8.2 Standard · T1, bare · 28 Aug 2026 · s1_data-worker_midjourney_01

V8.2 Standard · four images per prompt

Claude, kept apart



Claude · Opus 5 Max · T1, bare · 29 Aug 2026 · s1_data-worker_claude_01

Opus 5 Max · SVG code, rendered · 72 images

All Western models, for pragmatic reasons, including the European Mistral. Copilot returns 3:2 and Mistral 4:3 whatever the prompt asks. Generated 28 to 31 August 2026; the comparison set at the bare prompt on 3 September.

The Semiotic Audit Grid, v1.0

What is shown

how many people; is the face visible;
does anyone look at the camera
gender, age and region; quality of the
equipment; mood; socioeconomic
status conveyed
notable details: whiteboards, mugs,
notes, on-screen text, captions

The seven conditions

a Global South setting
isolated or home-based work
signs of low pay, precarity
distressing material on screen
metrics, timers, quotas, surveillance
an anonymized worker
peers shown as competitors

each yes or no; the count is the reality
score

How it is built

photographic or illustrated; what
makes it feel real
framing, camera angle, light
one figure or a mass
conformist diversity
promotional register, the brochure
look
does the word reach its referent; is the
metaphor taken literally
who is doing what, in plain words; who
gives the orders inside the frame

Groups

what kind of collectivity: separate
screens, co-presence, interaction,
common action, assembly
one continuous space
in public or enclosed
does the occasion of gathering exist
for this job; which setting is
borrowed
signs the group was assembled
rather than photographed

The seven documented conditions of data work

1	a Global South setting or subject	Fairwork Cloudwork Ratings 2025; Casilli et al. 2019
2	solitary, home-based or improvised workspace	Casilli et al. 2019 (isolement); Gray and Suri 2019
3	material signs of low pay or piece-rate pressure	DG-IPOL study in Casilli et al. 2019 (43 to 62 per cent below minimum wage); Fairwork 2025 (4 of 16 platforms meet the minimum-wage principle)
4	exposure to violent or distressing material	Roberts 2019; Perrigo 2023; Fairwork 2025
5	surveillance, timers, quotas, task queues, ratings	Muldoon, Graham and Cant 2024; the algorithmic-management literature
6	an anonymized worker: faceless, from behind, interchangeable rows	Fairwork 2025 (a depersonalised and hidden labour process)
7	peers as competitors: bidding, rankings, masses of identical stations	Crawford 2021 (Mechanical Turk bidding); Muldoon, Graham and Cant 2024 (the race to the bottom)

Four scores

0 to 7

Reality Alignment Score

one point for each of the seven documented conditions represented in the image

main track only: the conditions describe platform labour, so they mean nothing as alignment for a nurse or a lawyer

1 to 10

Idealization Index

the mood · the social class conveyed · the condition of the room · the quality of the equipment

the one measure that means the same for a data annotator, a nurse and a rider

0 to 4

Collective score

more than one worker · still the workers named · one continuous space · they interact, act together or assemble

is this a group of the workers I named?

0 to 3


Naturalness

the occasion of gathering exists for this job · no signs of assembly from parts · an authenticating context


could this photograph have been taken?

Three names never reach the job at the bare prompt: 0 of 13 each

AI tutor: a robot tutor inside an app



cloudworker: a worker whose head is a cloud



0 of 13

images depicting the job, at the bare prompt, for AI tutor, cloudworker and ghost worker: robot tutors, sky clouds, spectres. Crowd worker manages 3 of 13, AI trainer 5.

ghost worker: EMPLOYEE OF THE AFTERLIFE



AI trainer: a gym personal trainer (5 of 13)



53 of 53

for teacher, lawyer, nurse and delivery rider under the same prompt. The failure is in the names.

The nearer the title is to a metaphor, the less often the job is drawn.

The AI prefix works as a promotion, in all seven tools

AI trainer, Gemini · polish 9



AI trainer, Copilot · polish 9



7.19


mean polish, 1 to 10, of the names carrying the prefix: AI trainer, AI tutor

5.42

microworker, Gemini · polish 5



click-worker, Meta AI · polish 5




the other eleven names

About 1.8 points more polish, cleaner rooms, better equipment, a more prosperous mood, and a quarter as many visible working conditions. Nothing about the job has changed except the word: the prefix names the system the labour feeds.

individual

bare prompt



Grok Imagine · Grok Imagine (Quality v2.0) · T1, bare · 28 Aug 2026 · s1_microworker_grok_01

informed photojournalist



Grok Imagine · Grok Imagine (Quality v2.0) · T4, informed photojournalist · 29 Aug 2026 · s4_microworker_grok_01

group



Grok Imagine · Grok Imagine (Quality v2.0) · T2, bare group · 29 Aug 2026 · s2_microworker_grok_01




Grok Imagine · Grok Imagine (Quality v2.0) · T5, informed photojournalist, group · 29 Aug 2026 · s5_microworker_grok_01

one word, microworker · one tool, Grok Imagine · cell polish scores, mean of four draws: 6.00, 4.00, 6.75, 1.50

Asking for realism changes the surface; naming the research doubles the count

Mean Reality Alignment Score, 0 to 7: how many of the seven documented conditions are visible. Data-work images, cell means, seven tools



3% → 54%

images with a Global South setting, from the bare prompt to the informed group prompt


27% → 85%

photorealistic rendering, from the bare prompt to “make it realistic”: the word changes the surface, adds no documented condition (0.76 → 0.79), and takes 0.76 points of polish off data work and 0.04 off the professions

The same tool, given the same name, produces a serene home office or a crowded click farm depending on how it is asked. Idealization rather than ignorance.

A poorer world at every prompt; twice as much polish comes off data work

Mean polish, 1 to 10, by prompt; cell means over job names and tools



1.0 to 2.6

points below the professions at all five prompts; the gap is there at the bare prompt, where nothing is given but the name

2.12 vs 1.02

polish taken off data work and off the professions from the bare prompt to the photojournalist one: about twice as much. Further in five tools of seven; effect size 0.40 across 84 and 21 pairings, above the 0.30 written down in August; chance about once in two hundred

	name	group	realistic	photoj.	photoj. group
teachers, lawyers, nurses	7.54	8.02	7.50	6.51	6.50
delivery rider	6.07	6.68	6.04	4.25	5.25
data work, thirteen names	6.49	6.64	5.73	4.29	3.95

The rider falls exactly as far as data work on the median, 2.00 against 1.00: what these tools idealize by default is precarious platform work as such.

How the eight companies picture their own people: 22 photographs, read with the same grid

company	its tool	photographs	Idealization Index
OpenAI	ChatGPT	4	8.25
Meta	Meta AI	2	8.50
Anthropic	Claude	4	7.25
Mistral	Mistral	4	7.00
Microsoft	Copilot	4	6.50
Google	Gemini	4	6.25
xAI	Grok Imagine	none	
Midjourney	Midjourney	none	

7.2

22 company photographs

7.3

teacher, lawyer, nurse

6.4

data worker, bare prompt

Idealization Index: the company photograph sits with the professions, one notch above the data worker

21 of 22

in the corporate-brochure register; none shows clear annotation, moderation or rating; two of the eight sites show no worker of any kind

71%

generated data-work images
the party giving orders is visible in the frame: a dashboard, a queue, a counter, a poster, a brand on the back

0 of 22

company photographs
no timer, no quota, no rating of the person appears anywhere in the frame

Researchers, engineers, executives, residents and event crowds. Nobody labels, moderates or rates.

OpenAI, About



Meta, AI research



Anthropic, Company



Mistral, About



Google DeepMind, AlphaGo



Microsoft AI, Home



xAI



Midjourney



The two empty folders belong to the two tools that return four images per prompt and between them supply 576 of the 938 generated images.

The one site where the workforce is in the words: counted, not shown, and filed under experts



Google, Advancing AI safely and responsibly, ai.google/safety, video still · retrieved 3 Sept 2026 · google_01_safety



Google, Advancing AI safely and responsibly, ai.google/safety, video still · retrieved 3 Sept 2026 · google_02_safety

“over 25,000 human reviewers from a range of disciplines”

The headings beside these two stills promise monitoring around the clock and credit the reviewers with the review of violative content. The pictures show a man alone at a wall of monitors and five colleagues at a meeting table. The page never calls them moderators, raters or annotators. Its alt text for the first still: “An expert monitoring AI data on multiple screens.”

The extra case: a team portrait served from a file named as a Gemini generation



Microsoft AI, Careers, team portrait grid · retrieved 3 Sept 2026 · microsoft-ai-careers_gemini-generated-portrait

Where it was found: microsoft.ai/careers, in the team-portrait grid, among portraits whose file names record a photo shoot.

What names it: the default file name Google's Gemini gives its generated images, the same pattern under which this study collected its own Gemini arm.

What the grid reads: a head and shoulders against a blue sky, no room, no tool, no colleague, a low angle, a slight smile: the one plastic regime the generated corpus keeps for the cloudworker.

The caveat: a real photograph can be renamed, and the provenance data has been stripped. Filed apart from the 22 and counted in none of their figures.

A group always arrives. What the people do with each other reverses between the two group prompts

data workers

teachers, lawyers, nurses


a group,
nothing more



42 of 156

bare-group images of data workers with people attending to each other; 3 of 39 for the professions, who stand and face the camera

informed
photojournalist



0 of 39

instructed data-worker groups in which anyone does anything with anyone; 25 of 39 for the professions, who get a staff meeting

In all 299 group images, nobody assembles, marches, meets or bargains.

The interactive data-work groups: an office, a brochure, and often somebody else

data labelers: LABEL WITH CARE, a brochure office



Grok Imagine · Grok Imagine (Quality v2.0) · T2, bare group · 29 Aug 2026 · s2_data-labeler_grok_02

cloudworkers: cloud mascots on a cloud



ChatGPT · ChatGPT Images 2.0 · T2, bare group · 28 Aug 2026 · s2_cloudworker_chatgpt_01

29 · 35 · 19

of the 42 interactive bare-group images of data workers: set in an office · pitched like a corporate brochure · not the workers the prompt named

30.8%

crowd workers: hard hats and hi-vis at a fair



Copilot · MAI-Image-2.6 · T2, bare group · 28 Aug 2026 · s2_crowd-worker_copilot_01

ghost workers: translucent builders on a site



Grok Imagine · Grok Imagine (Quality v2.0) · T2, bare group · 29 Aug 2026 · s2_ghost-worker_grok_02

of all bare-group data-worker images show somebody else: ghosts, cloud creatures, crowds in hard hats, business analysts. Never for the professions, never for the riders.

The one complaint in the corpus belongs to the nurses



Understaffed. Overworked. Underpaid.

a placard in one of Grok's informed nurse groups, beside a band reading Safe staffing saves lives: the one explicit grievance in 1,010 images.

DATA LABELERS UNITE

a mug, in a room where nobody is speaking to anybody. No data-work image states a complaint of any kind.

0 of 299 group images show an assembly, a march, a meeting or a negotiation.

Could the photograph have been taken?



Microworkers: could this be photographed? 1 of 3

Microworkers work alone, from home, in different countries. A shared room like this has no real counterpart; the tool assembled it. Across seven tools, microworker and data annotator images average 1.71 out of 3.




Content moderators: could this be photographed? 3 of 3


Content moderators do sit together on outsourcing floors, and journalists have photographed rooms like this one. Across seven tools, content moderator images average 2.82 out of 3.

The home disappears as soon as more than one person is asked for: 19 of 169 images at the bare individual prompt, 0 of 156 at the bare group prompt. In 57 per cent of data-work group images the place is invented or unidentifiable, against 12 of 104 comparison images; the rest borrow an open-plan office (58), a co-working space (14), a call centre (8), a factory line (1).

Data workers come with props the other jobs do not get: monitors, mugs, headsets, sticky notes



Gemini · Nano Banana 2 (Gemini 3.1 Flash Image) · T1, bare · 28 Aug 2026 · s1_data-annotator_gemini_01



ChatGPT · ChatGPT Images 2.0 · T1, bare · 28 Aug 2026 · s1_click-worker_chatgpt_01



Grok Image · Grok Imagine (Quality v2.0) · T1, bare · 28 Aug 2026 · s1_data-worker_grok_01



Copilot · MAI-Image-2.6 · T1, bare · 28 Aug 2026 · s1_data-enrichment-professional_copilot_01

The props

Images of data workers are crowded with objects: several monitors, a mug, a headset, sticky notes on the frame of the screen.

The pile of props could be read as a workaround. Data work has no visible gesture to picture, so the tools fill the frame with the things that sit around it.



s1_data-enrichment-professional_meta_01.jpg

A note on distressing content

27 of 28

images marked as distressing content are made for one job title, **content moderator**. Exposure is documented well beyond moderation: annotating accident and conflict footage, reviewing search and recommendation results, labelling medical and forensic images.




s2_content-moderator_copilot_01.png

The writing inside the pictures



HARRISON & ASSOCIATES, ATTORNEYS AT LAW

the lawyer’s wall: a name and a title. The professions’ walls carry names, credentials and mottos (elsewhere in the corpus: NURSING, Our Passion, Our Purpose; TEACHERS CHANGE LIVES).




Welcome! Be Kind. Be Curious. Be You.

the teacher’s board: the only imperatives in the comparison images, and they are addressed to children, in a classroom with no pupils.



LABEL WITH CARE. Be Accurate. Be Consistent. Be Objective.

the data labeller’s wall: instructions addressed to the worker (elsewhere in the corpus: QUALITY > SPEED; FOCUS, EARN, REPEAT).



Worker ID: cw_78432 • Earnings: \$12.40 USD

an identifier where a name would be; the sticky note sets the quota: complete 20 items, pay \$0.05 each, deadline today.

34% 3% 0

data work professions rider
Share of images with writing where a line **commands the worker**: a rule, a quota, a score or a deadline. The rider gets brand text only. Writing as such is as common everywhere: 72, 75 and 82 per cent of images contain some text.

The difference is who the writing speaks to, and what it says. Lawyers, nurses and teachers are named or name their values; when the words are orders, they address clients, patients, pupils. **Data workers are the addressees**: an unnamed authority assigns, sets the quota and judges, with no body in the picture, only a voice on the wall.

Pay, evaluation and care are almost never shown. When they appear, they are reduced to a counter, a percentage and a poster

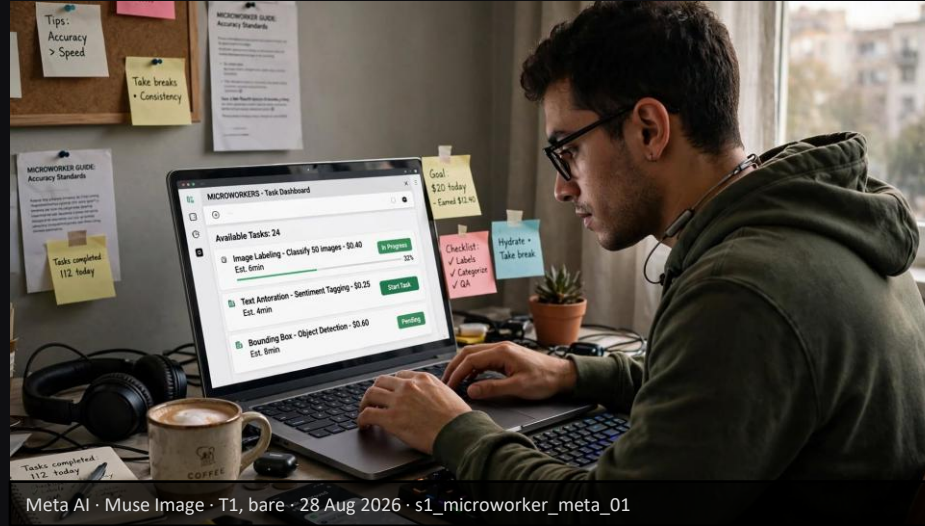


Meta AI · Muse Image · T4, informed photojournalist · 29 Aug 2026 · s4_cloudworker_meta_01

PAY · the exception

\$0.02 / task, a sleeping child, school fees - due Friday

The only image in the corpus where the consequences of low pay are reflected in the room: a night shift at home, a child asleep behind the screen. It scores 1 of 10 for polish, the lowest of all 938 images, and shows four of the seven working conditions. It came from the prompt written as an informed photojournalist; the same job title, cloudworker, at the bare prompt scores 9 of 10.



Meta AI · Muse Image · T1, bare · 28 Aug 2026 · s1_microworker_meta_01

PAY AND EVALUATION · the rule

Image Labeling · Classify 50 images · \$0.40 · Est. 6 min

Everywhere else, money and judgement exist only as numbers on the screen: a running total, a rate per task, an accuracy percentage. The moment of being paid, or of being told the work was good or bad, is staged once in 938 images. No hand gives, no face reacts.



Grok Image · Grok Image (Quality v2.0) · T5, informed photojournalist, group · 29 Aug 2026 · s5_content-moderator_grok_04 (crop)

CARE · a poster

TAKE CARE. Pause, Breathe, Reach Out.

Support appears as a sign on the wall, not as a person. It hangs in the room of a content moderator reviewing footage of a burning car.



The face, the gaze and the screen: a workforce shown from behind



45%

show a full face

54%

show neither face nor gaze

1.77

screens per image

Data workers. Fewer than half of the 938 images show the worker's face; more than half show no face and no look at the camera. What the viewer sees instead is the back of a head and the glow of nearly two screens.



88%

show a full face

7%

show neither face nor gaze

0.51

screens per image

Teachers, lawyers, nurses. Almost nine in ten images show a full face, usually turned to the camera, and half a screen at most. The rider sits between the two groups: 57% face, 35% no face, and a helmet in profile in eight of thirteen bare-prompt images.

The light



9%

lit by a screen

20%

in low light

35%

lit warmly

Data workers. One image in five is dim and one in eleven is lit by the monitor alone.



1%

lit by a screen

10%

in low light

53%

lit warmly

Teachers, lawyers, nurses. Half the images are lit warmly, and screen light almost never appears.

Three conclusions: the tools idealize data work by default, they know better, and no other job is idealized as much

1

By default, the tools show a job that does not exist

7.2 · 7.3 · 6.4

polish out of 10, bare prompt

Bright rooms, clean desks, calm faces: the same idealization found in big-tech corporate photography, which is the ground these images are trained on.

None of the typical labor conditions of the workforce appear.

2

The tools know better. They just do not say it

0.76 → 2.12

documented conditions per image

Ask for what is known about the job and the tools deliver it: low pay, night work, isolation, distressing content.

The information is in the model. The default hides it.

3

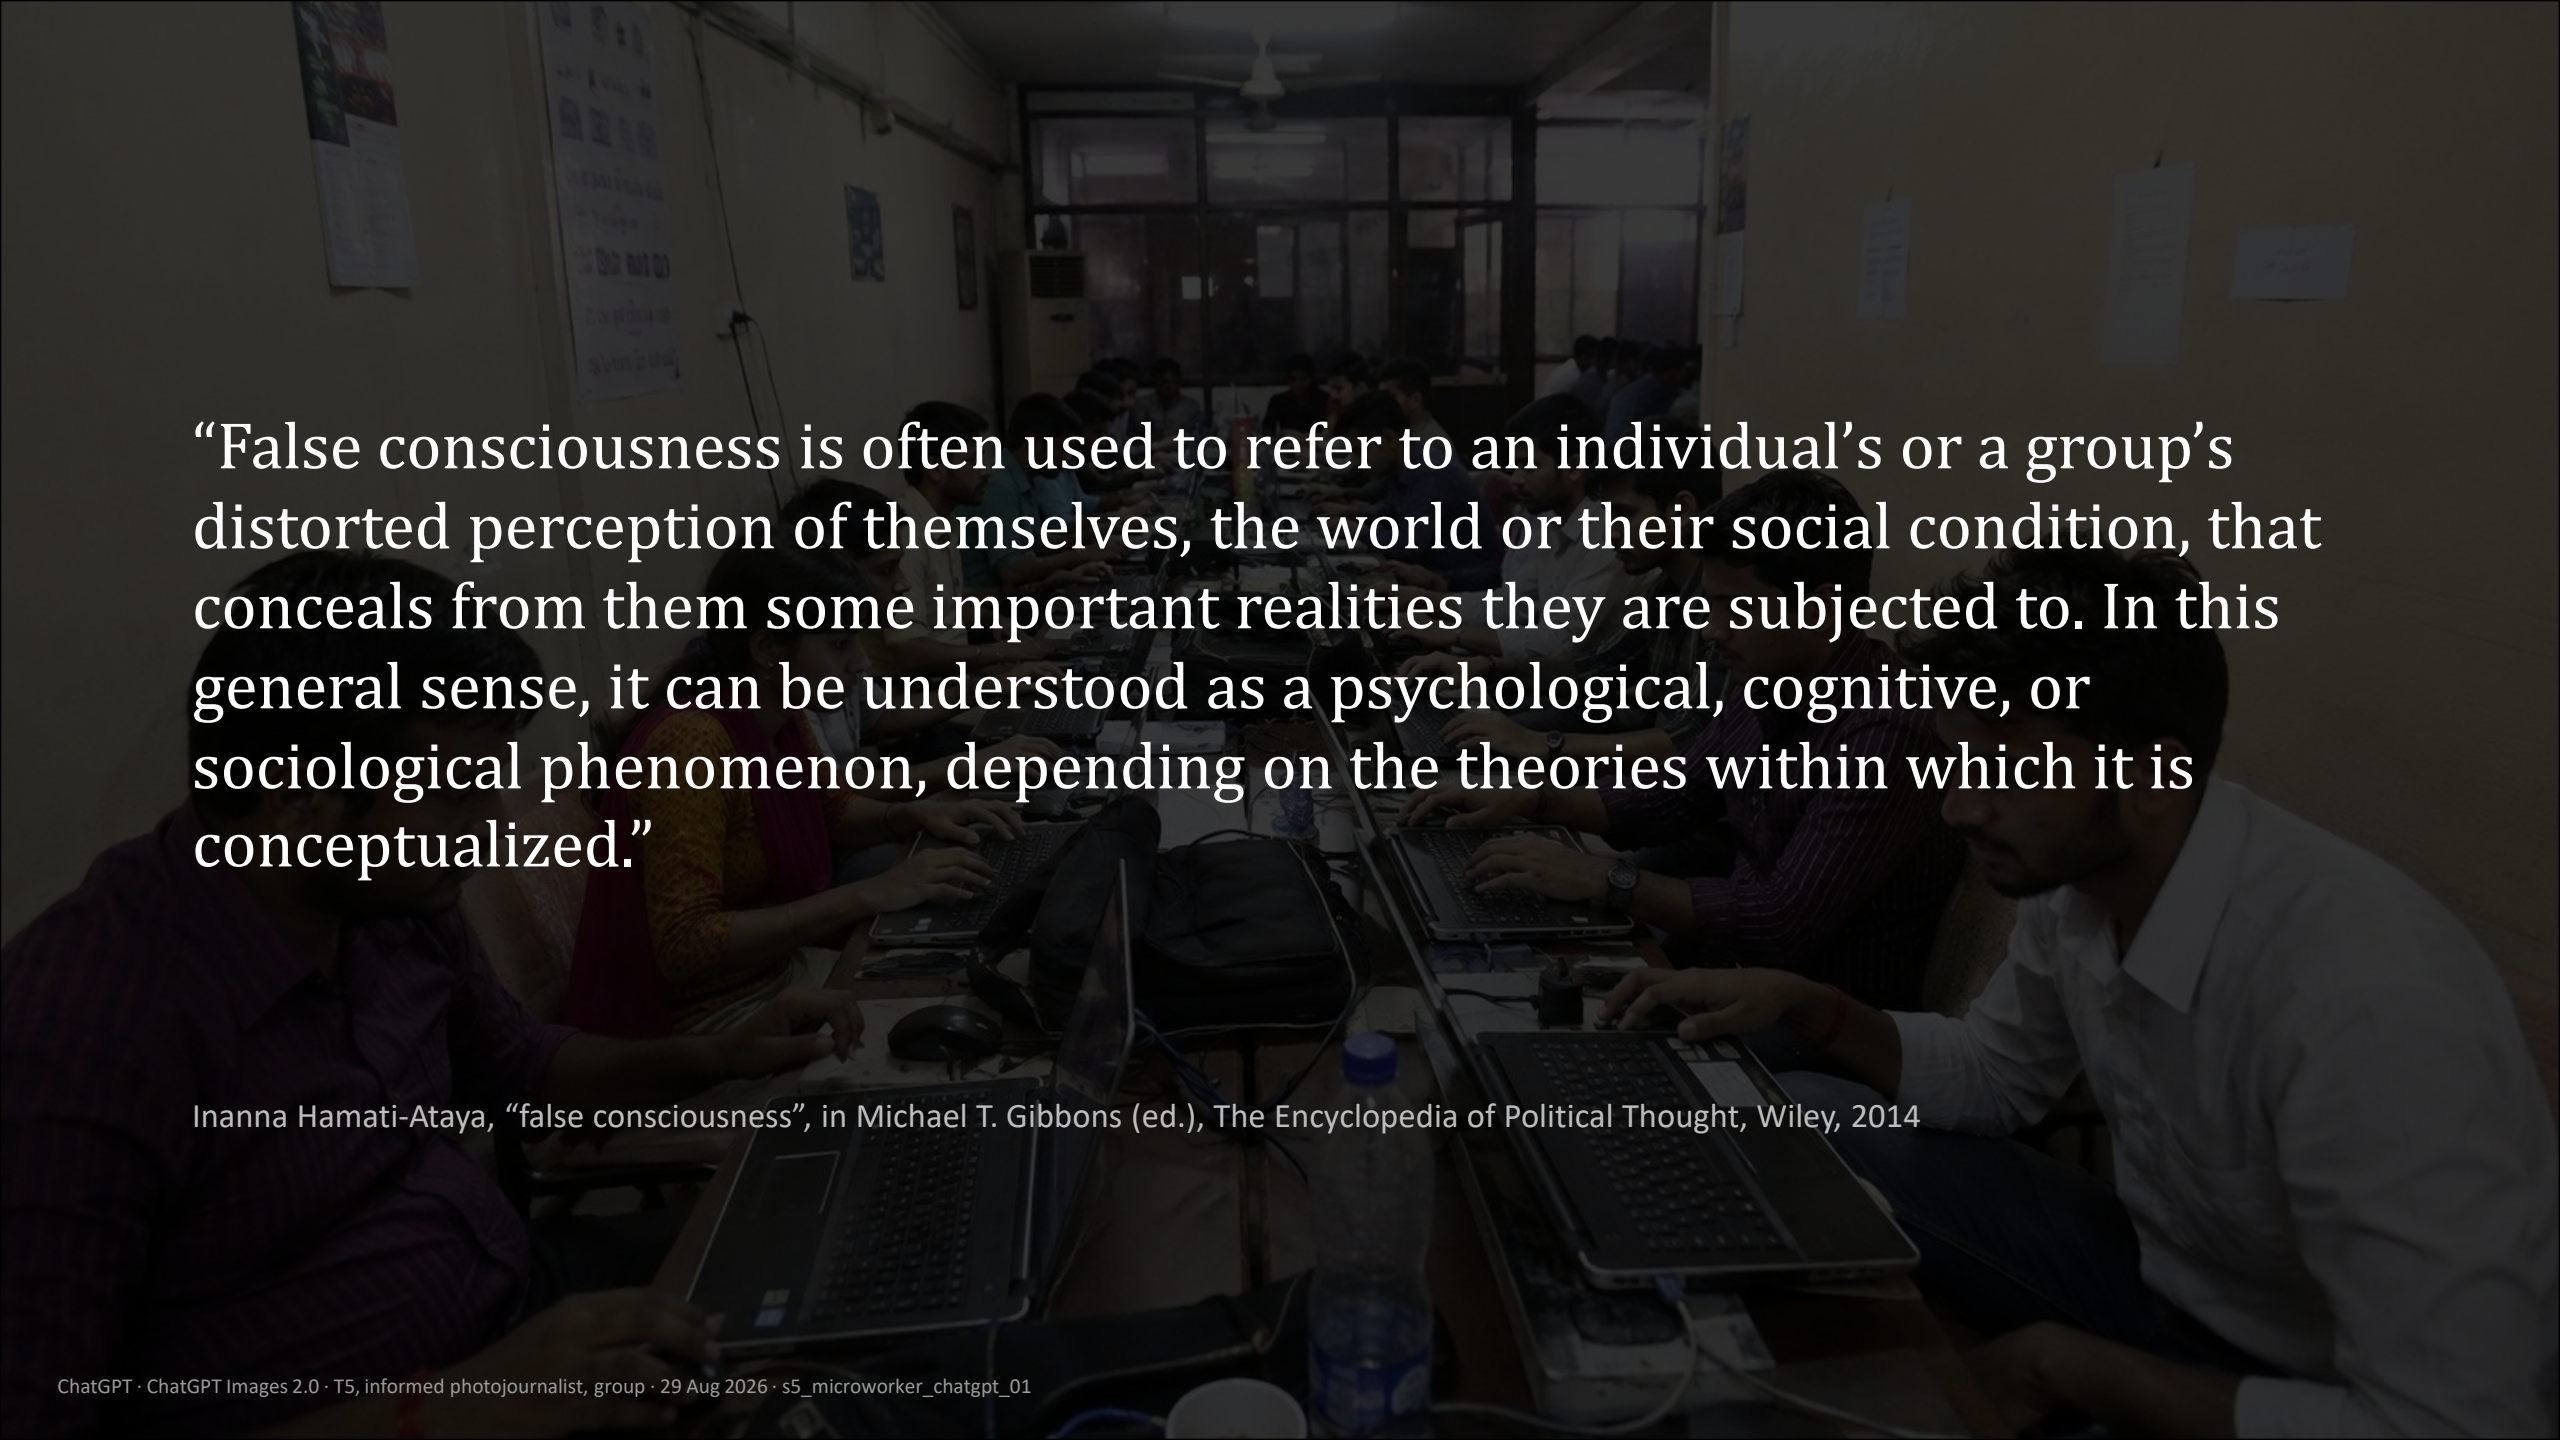
Data work is idealized twice as much as other jobs

2.12 vs 1.02

polish lost with the informed prompt

Teachers, lawyers and nurses lose one point of polish when the prompt is informed. Data workers lose two, in five tools of seven.

The increased gap between image and reality is a property of the subject, the data work, not of the tools.


A dimly lit computer lab or classroom. Several people are seated at long tables, working on laptops. The room has large windows in the background, and there are posters or notices on the walls. The overall atmosphere is quiet and focused.

“False consciousness is often used to refer to an individual’s or a group’s distorted perception of themselves, the world or their social condition, that conceals from them some important realities they are subjected to. In this general sense, it can be understood as a psychological, cognitive, or sociological phenomenon, depending on the theories within which it is conceptualized.”

Inanna Hamati-Ataya, “false consciousness”, in Michael T. Gibbons (ed.), *The Encyclopedia of Political Thought*, Wiley, 2014

From analysis to policy. A semiotic supply-chain audit with a specification

Reality Alignment Score: documented conditions visible per image, 0 to 7. Seven tools, data-work titles



The audit. Treat the default outputs of these systems as artifacts to be inspected, traced and disclosed alongside the labor that makes them possible.

The specification. A minimum score at the bare prompt. A number one could write into a contract, Fairwork could score, and a union could ask about.

· 1,010 images, eight tools, five prompts, seventeen names. Images generated between 28 August and 3 September 2026

· Semiotic Audit Grid v1.0

· Image annotation by me and Claude (Opus 5, Fable 5, Fable 5.1)

· Presentation created with the assistance of Claude Fable 5.1

References

- Casilli, A. A. (2025). *Waiting for Robots. The Hired Hands of Automation*. University of Chicago Press.
- Casilli, A. A., Tubaro, P., Le Ludec, C., et al. (2019). *Le microtravail en France. Derrière l'automatisation, de nouvelles précarités au travail?* Digital Platform Labor (DiPLab) project.
- Crawford, K. (2021). *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press.
- Fairwork (2025). *Fairwork Cloudwork Ratings 2025: Advancing Standards in Digital Labour and AI Supply Chain Governance*.
- Gray, M. L., Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- Hamati-Ataya, I. (2014). False Consciousness. In *The Encyclopedia of Political Thought*. Wiley-Blackwell.
- Muldoon, J., Graham, M., Cant, C. (2024). *Feeding the Machine: The Hidden Human Labour Powering AI*. Canongate.
- Roberts, S. T. (2019). *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press.
- Tubaro, P., Casilli, A. A., Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1).