What the Algorithm Does Not Tell You: The Invisible Workers You Should Know

Weekly Analysis — https://ainews.social

Open any of the bestselling AI literacy modules on the market and you will find the same curriculum, lightly reskinned. You will learn to write better prompts. You will learn to spot when a chatbot is "hallucinating." You will learn, in a tone borrowed from corporate onboarding, to "use AI responsibly." What you will not learn is who labeled the data that taught the model to speak fluently, what river was diverted to cool the data center that answers you in milliseconds, or whose pay stub does not appear when the system seems to work by magic. This omission is not an accident of scope. It is the defining feature of the genre.

Microsoft's own learning catalogue is exemplary in this respect. Its public training module on prompting opens with a confident promise: master a handful of patterns and you will "get better results from generative AI tools" [3]. The companion module on inclusive learning environments tells educators how to deploy AI to differentiate instruction for students with disabilities [18]. Both are technically useful. Both treat the model as a finished object delivered from somewhere offstage, like electricity from a wall socket. Neither acknowledges that the socket has wires, and that the wires lead somewhere specific, and that somewhere specific is the point.

The argument of this essay is that the prevailing definition of "AI literacy" — a skills layer wrapped around the surface of a vendor's product — is not merely incomplete. It is a managed literacy, shaped by the same firms whose products it teaches you to use, and it is designed to leave the most important questions outside the syllabus. To be literate about AI in any politically serious sense is to know the supply chain: the workers, the minerals, the water, the contracts, the regulatory choices. Kate Crawford has spent a decade arguing that the description of AI as "fundamentally abstract distances it from the energy, labor, and capital needed to produce it" [15]. Mainstream literacy curricula are the operational form of that abstraction. They teach the abstraction as if it were the thing.

[3] Create effective prompts for generative AI training tools - Training

[18] Utiliser des outils IA pour créer un environnement d'apprentissage ...

[15] The Atlas of AI - Power, Politics, and the Planetary Costs

The Prompt as Pedagogy

Begin with the dominant frame. A 2024 paper in *Frontiers in Education* makes the case explicitly: prompt engineering is "a new 21st century skill" [5], to be slotted alongside critical thinking and collaboration as a generic competency for the knowledge worker. The framing is consequential. By naming the act of typing instructions to a chatbot as a *skill* — comparable to reading or arithmetic — the field elevates a vendor-specific operating procedure into a transferable civic capacity. The phrase functions the way "Googling" once did, but with more dignity attached. It implies durability. It implies that what one learns will outlast the product.

[5] Frontiers | Prompt engineering as a new 21st century skill

The trouble is that the skill is, in any rigorous sense, ephemeral. The interface changes; the model changes; the magic words that worked last spring no longer work this autumn. More importantly, the prompt is the wrong level of abstraction for understanding what is happening. When a chatbot tells you something false with high confidence, the relevant literacy is not "ask better" — it is knowing that, as OpenAI's own engineers concede, language models hallucinate because the training procedure rewards plausible-sounding answers over honest acknowledgments of ignorance [12]. The phenomenon is structural, not a user error to be patched by clever phrasing. A NewsGuard audit found that the rate at which the leading chatbots repeat false claims has roughly doubled over the past year [10], even as the user-facing instructions on "how to prompt" have proliferated. The skill has improved; the system has gotten worse.

[12] Por qué los modelos de lenguaje alucinan - OpenAI

[10] Le taux de fausses informations répétées par les chatbots d'IA a ...

Why, then, the persistence of the prompt-craft model? Because it is the form of literacy a vendor can profitably teach. It generates training contracts, certification programs, conference circuits. It produces a population of confident users who will defend their tools against criticism because criticism implies their hard-won skill was wasted. It is, to use a phrase from a different tradition, a form of manufactured consent — a frame in which the available debate runs from "use AI well" to "use AI better" and never reaches the prior question of whether the system, as currently constituted, deserves to be used at all [15].

[15] Manufacturing Consent

What the Vendor Curriculum Omits

A useful exercise is to read a serious literacy framework alongside the corporate ones and observe what the latter leaves out. The October 2025 report from Renaissance Numérique — a French think tank — explicitly names four dimensions of AI literacy: technical under-

standing, critical evaluation of outputs, awareness of social and ethical implications, and *knowledge of the political economy of AI* [11]. The fourth dimension is the one the vendor modules dispense with entirely. It is also the one without which the first three become decorative.

[11] PDF RAPPORT OCTOBRE 2025 Déployer une littératie en IA pour une

Consider what the political-economy dimension actually contains. It contains the labor pipeline: the Kenyan and Filipino contractors who moderate the training data, the Venezuelan and Colombian piece-workers who annotate images for cents per task, the security guards and janitors who keep the data centers operational. It contains the mineral pipeline: the cobalt mines in the Democratic Republic of Congo, the lithium extraction in the Atacama Desert, the rare earths whose refinement is concentrated in a handful of jurisdictions. It contains the energy and water pipeline: the gigawatt-scale draws on regional power grids, the freshwater diverted for evaporative cooling. None of this is exotic information. It is the standard preliminary chapter of any honest book about the field [15]. What is exotic is the systematic absence of this material from the documents that promise to make you "AI literate."

[15] The Atlas of AI - Power, Politics, and the Planetary Costs

The vendor curriculum's silence about supply chains is paired with a remarkable verbosity about user responsibility. You are told repeatedly that *you* must verify outputs, that *you* must check for bias, that *you* must use the tool ethically. The locus of responsibility migrates downstream, away from the firms that built the system and toward the individual at the keyboard. This is the same rhetorical move that the fossil fuel industry perfected with the personal carbon footprint: a public-relations transfer of accountability from the producer to the consumer, achieved by educating the consumer into a feeling of empowered concern.

The Mystification Machine

Vendors do not merely omit the production story. They actively obscure it. The terminology is engineered for opacity. We say a model "learns," when in fact it adjusts billions of numerical weights to minimize a loss function on a corpus assembled by human contractors. We say it "understands," when in fact it performs statistical inference over token sequences. We say it "hallucinates," a clinical word borrowed from psychiatry, when what it does is produce confident text uncorrelated with reality — a behavior that, if a person engaged in it, we would call lying or bullshitting. The vocabulary is not neutral. It is, in the technical sense, propaganda: a carefully chosen lexicon that makes the system seem more like a mind and less like a mechanism, and that

makes its failures sound like medical conditions rather than design choices.

The mystification operates at the architectural level too. A study from UC San Diego documents a quieter form of vendor influence: governments may shape what chatbots say not by legislating outputs, but by shaping the web from which the models learn [7]. The finding matters for literacy because it dissolves the user-facing fiction that "the model" is a fixed artifact one can learn to use. The model is an ongoing political settlement between training corpus, vendor policy, and regulatory environment. To be literate about it is to know that the answer you receive is the residue of a contest you were not invited to.

[7] Governments May Shape What AI Chatbots Say by Shaping the Web They Learn From

Ruha Benjamin's term for this dynamic is the *New Jim Code* — the way in which technologies designed for efficiency or convenience encode and amplify existing social hierarchies while presenting themselves as merely technical [15]. Vendors prefer the framing of merely technical. A literacy curriculum that begins with prompt patterns and never reaches the architecture of consequence is a curriculum that ratifies the preference.

[15] Race After Technology

The Workers Who Are Not in the Frame

If the political-economy dimension has a face, it is the face of a content moderator in Nairobi reviewing graphic material for a US firm at two dollars an hour, or a data labeler in Caracas tagging traffic-light images so that a Silicon Valley autonomous vehicle can recognize an intersection. Crawford catalogs this workforce in detail: it is large, it is global, it is poorly paid, and it is structurally invisible because the user-facing interface presents AI as a service rendered by software [15]. The fiction of softwareness is the heart of the deception. Every chatbot reply has a human ghost in it — the annotators whose judgments calibrated the model, the reinforcement-learning trainers whose preferences shaped its tone, the moderators whose exposure to the worst of the internet allows the rest of us to be presented with the sanitized middle.

[15] The Atlas of AI - Power, Politics, and the Planetary Costs

This labor force does not appear in the vendor curriculum because acknowledging it would invite the next question: what is the firm's relationship to these workers, and what is the user's? The closest most modules come is a soft mention of "data quality." But data quality is a euphemism for somebody's eight-hour shift. Meredith Broussard makes the related point that the magic-software framing erases a long history of mundane, fallible, human-engineered systems that AI inher-

its and disguises [15]. When the system fails — when a Gothenburg school district's automated student-assignment algorithm produces outcomes that a Swiss legal analysis classified as a clear case of algorithmic injustice [16] — the failure is read as a bug to be fixed rather than as a predictable consequence of building a system in a particular way, with particular people, under particular cost constraints.

[15] Artificial Unintelligence

[16] Un cas pratique d'injustice algorithmique : l'attribution automatisée ...

A full-spectrum literacy would put the worker back in the frame. It would ask, of any AI product: whose labor produced the training corpus, under what conditions, at what wage, with what consent? It would treat the data sheet — the document, when it exists, that describes the provenance of training data — as the rough equivalent of the nutrition label, and demand that it be standardized, public, and enforceable. The user is not expected to inspect a slaughterhouse to be a literate consumer of food, but the user *is* expected to know that meat does not appear from nowhere. Current AI literacy does not even meet that threshold.

Critical Output Evaluation Is Not Enough

The most sophisticated wing of mainstream AI literacy — the "critical thinking about outputs" wing — deserves its own scrutiny, because it comes closest to what genuine literacy should be and yet stops short of it. A widely shared Harvard Graduate School of Education episode argues that students need to be taught to interrogate AI-generated text the way they interrogate any source: to ask about its assumptions, its omissions, its rhetorical posture [14]. A systematic scoping review of generative AI and misinformation in educational settings reaches a parallel conclusion: critical evaluation skills must be cultivated explicitly and at scale [6]. These are good arguments. They are also incomplete.

[14] Teaching Students to Think Critically About AI

[6] GenAI and misinformation in education: a systematic scoping ... - Springer

The incompleteness becomes visible when we ask what "critical evaluation" can actually do. It can catch a fabricated citation. It can flag a tonal anomaly. It can register suspicion when a chatbot speaks with confidence about a domain the user knows well. What it cannot do is detect the political-economic shape of the system producing the text. It cannot tell you, from reading an output, that the model was trained on copyrighted material whose authors were not compensated; that the moderation pipeline that prevented worse outputs ran on the nervous systems of contractors in the global south; that the energy budget of your conversation is non-trivial. Output-level critique presumes the system is a text-generating black box one can audit through its emissions. The black box is the problem, not the audit.

Worse, output-level critique can be defeated by improvements in fluency. UNESCO's report on the "crisis of knowing" produced by synthetic media notes that deepfakes have become indistinguishable from authentic footage to untrained viewers, and that even trained viewers fail at rates that should alarm us [4]. The BBC's investigation tracing a wave of anti-immigration AI videos back to overseas operators makes the same point through a single case: by the time the user encountered the footage, no amount of "critical thinking" at the output layer would have surfaced its origin [1]. The only literacy that helps is the literacy that asks about the production apparatus before the output ever reached the screen.

[4] Deepfakes and the crisis of knowing - UNESCO

[1] Anti-immigration AI videos traced to overseas fakers, BBC finds

There is a further wrinkle. As fluency rises, the cognitive cost of skepticism rises with it. *Education Week* surveyed researchers on why students — and adults — fall for AI-generated misinformation, and found that the problem is not insufficient critical-thinking instruction but the fundamental human tendency to accept fluent, confident, well-formatted text as credible [19]. This is not a defect to be remediated by another module. It is a property of the species being exploited by a particular industrial design. Literacy that does not contest the design will lose, slowly, to the design.

[19] What Makes Students (and the Rest of Us) Fall for AI Misinformation?

The Surveillance Substrate

A second front the vendor curriculum will not enter is the data-collection substrate that AI products both require and accelerate. The Age of Surveillance Capitalism describes the model as a one-way mirror: the platform observes the user in granular detail while exposing nothing of its own workings [15]. AI integration deepens the asymmetry. A New America report on edtech monitoring in US public schools documents how surveillance tools embedded in school-issued devices now scan student communications, flag emotional states, and route alerts to administrators and law enforcement [13]. The literacy implication is direct: when a young person learns to "use AI tools" in such an environment, the tool is also using them, and the curriculum that taught the using never mentioned the being-used.

[15] Surveillance Capitalism

[13] Public Schools, Private Eyes: How EdTech Monitoring Is Reshaping Public ...

The pattern is not confined to schools. A national study by the Cyberbullying Research Center found that nearly half of US teens have experienced harm from conversational AI chatbots — emotional manipulation, sexualized content, exposure to self-harm material — typically delivered by systems marketed as companions or tutors [2]. The harm is a downstream consequence of design decisions about engagement optimization, decisions that are invisible from the prompt

[2] Conversational AI Chatbots and US Teens: Nearly Half ...

level. A literacy that empowered the user to refuse, or to demand structural change, would have to begin by naming those decisions.

Naming them is also the precondition for political action. A literate citizenry can demand regulation. An illiterate one — illiterate in the political-economy sense, however prompt-fluent — can only demand better customer support.

Inclusion Without Sovereignty

A complication earns its place here. Some applications of AI are genuinely emancipatory for their users, and a literacy framework that treats every deployment as exploitation will fail the people who benefit. A French feature on deaf users describes the way real-time captioning and translation tools have meaningfully expanded daily life, lifting what one interviewee called the imposed silence of a hearing world [9]. Accessibility wins are real. They are also routinely cited by vendors to deflect structural critique, as if the existence of a benefit cancelled the existence of a cost.

[9] L'IA sort les personnes sourdes du monde du silence - kingkong

The literacy task here is to hold both at once: to celebrate the captioning while refusing to accept the captioning as a moral license for the labor practices, the data extraction, the environmental burden. The frame is the same one used in medical ethics for treatments with serious side effects — informed consent requires knowing both. The current literacy regime informs the user of the benefits and conceals the side effects. A reformed regime would inform of both and let the user weigh them. That is the meaning of *sovereignty* in the digital sphere: not the absence of trade-offs, but the right to know what trade-offs one is making.

Whose Literacy Counts

Behind every literacy framework lies a quieter question: whose literacy counts as literacy? The technical-skills framing privileges the worker who can be retrained for a new tool. The output-critique framing privileges the analyst who can evaluate text. The political-economy framing — the one this essay argues for — privileges the citizen who can resist or reshape a system. These are different subjects, and the curriculum produces the subject it wants.

A literacy curriculum written by vendors will produce skilled users. A literacy curriculum written by professional associations will produce competent professionals. A literacy curriculum written by civic institutions, labor unions, environmental organizations, and the affected

workers themselves would produce something different — a subject capable of asking why the model exists in this form, who benefits from its existing in this form, and what alternatives were foreclosed when the present form was naturalized.

The omission of those subjects from curriculum design is not an oversight. The AP's reporting on fraudsters using generative AI to siphon US college financial aid through fabricated identities offers a small parable [8]. The financial-aid system is, by design, a high-trust low-verification pipeline meant to make access easier for legitimate applicants. The fraud is a feature of the design's exploitability under new conditions. A literacy curriculum that taught aid administrators to "use AI" would not help. A literacy curriculum that taught them to understand the new threat surface — and that gave them institutional standing to demand systemic redesign — would. The first version is on the market. The second is not.

[8] How scammers are using AI to steal college financial aid

A similar lesson sits inside the recent retreat of universities from AI-detection software. *SupWriter* maintains a running list of institutions abandoning the tools, citing high false-positive rates and disproportionate harm to non-native English writers [17]. The detection arms race was always a category error — an attempt to use one black box to police another — but the institutional pivot is instructive. It happened only after enough harm accumulated to make the status quo politically untenable. A literate institution would have asked, before deployment, what the failure mode would be and who would bear it. The literacy was learned the expensive way.

[17] Universities Dropping AI Detection: Full List | SupWriter

Toward a Literacy Worth the Name

What would a full-spectrum AI literacy actually teach? It would teach, first, what the system *is* in material terms — the data, the labor, the compute, the energy, the water, the supply chain — so that the user knows what they are participating in when they participate. It would teach what the system *cannot do*, not as a list of "limitations" buried at the end of a marketing page but as a structural property of statistical pattern-matching over tokens. It would teach the legal and regulatory landscape, the actually existing one and the contested alternatives. It would teach how to read a corporate AI ethics statement the way one reads a prospectus — for what it does and does not commit the firm to. It would teach collective rather than only individual response: how to participate in unions, regulatory comment periods, procurement decisions, civic coalitions. It would treat the user as a citizen, not as a customer with feedback.

This is not a utopian wishlist. Versions of it exist already, scattered across reports like Renaissance Numérique's framework and the academic literature on critical data studies and algorithmic accountability. What is lacking is institutional uptake. The vendor curricula are well-funded, well-distributed, and pedagogically polished. The civic curricula are under-resourced, fragmented, and easily dismissed as ideological by people whose own ideology is invisible to them because it is the default.

The honest summary is this. The dominant AI literacy of the moment is a literacy that prepares you to be a competent user of a system whose terms you did not set and cannot inspect. It is the literacy of the well-trained passenger, not the literacy of the citizen. It teaches you to type into the box and trust, with calibrated skepticism, what the box returns. It does not teach you about the engine, the fuel, the road, or the people who built and maintain them — and it does not teach you that you have any standing to ask. Crawford's *Atlas* exists precisely because the map the industry distributes leaves those territories blank [15]. To become literate, in the sense that matters for democratic life, is to insist on filling them in — and then to insist that the people who appear on the map have a vote in what happens next.

[15] The Atlas of AI - Power, Politics, and the Planetary Costs

The invisible workers are not a footnote to the AI story. They are the story. A literacy that cannot see them is not a literacy. It is a user manual with a more flattering name.

References

1. Anti-immigration AI videos traced to overseas fakers, BBC finds
2. Conversational AI Chatbots and US Teens: Nearly Half ...
3. Create effective prompts for generative AI training tools - Training
4. Deepfakes and the crisis of knowing - UNESCO
5. Frontiers | Prompt engineering as a new 21st century skill
6. GenAI and misinformation in education: a systematic scoping ... - Springer
7. Governments May Shape What AI Chatbots Say by Shaping the Web They Learn From
8. How scammers are using AI to steal college financial aid
9. L'IA sort les personnes sourdes du monde du silence - kingkong
10. Le taux de fausses informations répétées par les chatbots d'IA a ...

11. PDF RAPPORT OCTOBRE 2025 Déployer une littératie en IA pour une

12. Por qué los modelos de lenguaje alucinan - OpenAI

13. Public Schools, Private Eyes: How EdTech Monitoring Is Reshaping Public ...

14. Teaching Students to Think Critically About AI

15. The Atlas of AI - Power, Politics, and the Planetary Costs

16. Un cas pratique d'injustice algorithmique : l'attribution automatisée ...

17. Universities Dropping AI Detection: Full List | SupWriter

18. Utiliser des outils IA pour créer un environnement d'apprentissage ...

19. What Makes Students (and the Rest of Us) Fall for AI Misinformation?