The Literacy Divide: Who Learns to Command the Oracle?

Weekly Analysis — https://ainews.social

There is an oracle in the room, and a small priesthood has learned to address it. The rest of us — the overwhelming majority, including most of the people who will be governed, hired, diagnosed, sentenced, and graded by its outputs — are being trained to receive its pronouncements as weather. This is the shape of AI literacy in late 2026: a few learn to command, many learn to comply, and the gap between those two postures is hardening into a class line.

The word ”literacy” is doing enormous work in this discourse, and most of that work is invisible. When a Fortune 500 company says it has trained 80% of its workforce in ”AI literacy,” it usually means prompt scaffolding and tool adoption. When a school district says the same, it often means a vendor-supplied module on responsible use. When UNESCO says it, the term tries to carry the weight of civic judgment, source criticism, and a working theory of how these systems actually function. These are not the same thing, and the conflation is not accidental. The largest empirical study of undergraduate AI use to date, run out of Berkeley, found that access to AI tools and the sophistication of their use track existing class and institutional advantages so closely that the report’s authors describe a ”second digital divide” already entrenching itself inside the population that supposedly has access [25].

[25] The largest study of AI use by undergrads is in, revealing disparities in access — and in cheating

That second divide is the one this essay wants to map. Not the gap between those who have a ChatGPT subscription and those who don’t — that gap is real but trivially fixable — but the gap between those who understand what the system is doing when it answers them and those who don’t. The French government’s recent literacy report names this directly, distinguishing ”usage” from ”compréhension critique” and warning that conflating the two produces a population that is fluent at querying and illiterate at interrogating [6]. The conflation is convenient for vendors. It is catastrophic for everyone else.

[6] Déployer une littératie en IA pour une

Three Definitions Fighting for the Same Word

Walk through any week of AI-literacy discourse and you will find three definitions wearing the same coat. The first is operational: literacy as the ability to get useful work out of a model. This is the prompt-engineering definition, and it has the advantage of being measurable. You can test whether someone can decompose a task, chain a few prompts, evaluate an output for obvious error. A growing scholarly literature treats this as the foundational twenty-first-century skill, the new keyboard proficiency [22]. Systematic reviews of prompt-engineering curricula now read like the early 2000s spreadsheet-training literature: confident, scalable, and almost entirely silent on the question of what the tool is for [23].

[22] Prompt engineering as a new 21st century skill

[23] Prompt engineering in higher education: a systematic review to help

The second definition is comprehensional: literacy as understanding what the system is, how it was built, what data it was trained on, what its failure modes look like. This is the definition UNESCO has been quietly pushing through its competency frameworks, which insist that recognizing an AI artifact, understanding the difference between statistical and human reasoning, and being able to articulate a system's economic underpinnings are all prerequisites to anything that deserves to be called literacy [27]. It is also the definition that survives contact with hallucination. A reader who understands that a large language model is a stochastic next-token predictor with no access to ground truth is in a fundamentally different relationship to its output than one who has been told it is "intelligent." Harvard's misinformation researchers have proposed a whole conceptual framework around this distinction, arguing that the term "hallucination" itself is misleading enough to constitute a literacy problem of its own [18].

[27] AI competency framework for teachers - UNESCO

[18] New sources of inaccuracy? A conceptual framework for studying AI

The third definition is civic: literacy as the capacity to participate as an equal in decisions about how these systems are deployed in public life. This is the rarest framing and the most demanding. It requires the operational and comprehensional skills, plus a working grasp of the political economy in which AI sits — who owns the models, who labels the data, who absorbs the externalities. Carnegie's recent mapping of AI and democracy treats this kind of literacy not as a personal skill but as a precondition for legitimate governance, on the grounds that a citizenry which cannot evaluate AI claims cannot meaningfully consent to AI policy [2].

[2] AI and Democracy: Mapping the Intersections - Carnegie Endowment for

These three definitions are not nested neatly. They compete. A workforce trained to the first definition and credentialed in it is a workforce whose employers have every reason to suppress the second and third. The Berkeley study makes the trade visible in miniature: students who came in with what the researchers call "critical priors"

used AI less and questioned it more, while students who came in with operational fluency used it more and questioned it less [25]. The market rewards the second cohort. The republic needs the first.

[25] The largest study of AI use by undergrads is in, revealing disparities in access — and in cheating

The Vendor Curriculum

If you want to know what a society's definition of literacy actually is, look at who writes the curriculum. In AI literacy, that is increasingly the companies whose products the curriculum teaches. Microsoft, Google, OpenAI, and IBM have each rolled out free or near-free "literacy" tracks aimed at K-12 systems, community colleges, and workforce-development agencies. IBM's recent security portfolio launch was paired with what the company explicitly called "AI literacy uplift" for its enterprise customers, the literacy in question being almost entirely the operational kind — how to deploy, configure, and trust the vendor's tools [15]. Microsoft's accessibility training for educators teaches genuine techniques, but its frame is consistently "personalize learning using AI" rather than "evaluate whether AI personalization is what this learner needs" [21].

[15] IBM Brings Its Most Advanced AI-Powered Security Portfolio to Clients, and is Strengthened by Ongoing Project Glasswing Work

[21] Personnaliser l'apprentissage pour les étudiants handicapés à l'aide de

This is not a conspiracy. It is a structural fact about who has the resources and the motive to produce literacy materials at scale. But it has a consequence the field rarely names: the operational definition of literacy is being normalized as the whole of literacy, because the actors who can afford to teach are the actors whose interest in critical scrutiny is, charitably, modest. Zuboff identified the underlying logic years ago when she observed that surveillance capitalism's preferred response to skill displacement is "urgent IT upskilling" rather than any reconsideration of whether the displacement was desirable in the first place [27]. Upskilling defines the problem as worker readiness. It cannot define the problem as system design, because it is produced by the system designers.

[27] The Age of Surveillance Capitalism

Common Sense Media's survey of how AI has actually entered American children's lives makes the asymmetry concrete. Children are using these tools heavily, often without supervision, and almost always through interfaces designed by the companies whose business model depends on continued use [1]. The "literacy" content these children receive, when they receive any, is overwhelmingly produced by the same companies. A generation is learning what AI is from the people selling it. This is roughly equivalent to teaching nutrition exclusively through materials produced by the cereal industry, and the field has somehow convinced itself that this is acceptable because the cereal companies are also funding the schools.

[1] 2024 The Dawn of the AI Era - Common Sense Media

What "Critical" Actually Demands

The shorthand "critical AI literacy" gets thrown around as if its meaning were obvious. It isn't. To be critically literate about AI is to hold simultaneously a working understanding of how the system functions technically, how it was produced economically, where its outputs fail epistemically, and how its deployment reshapes the social field. Each of these is a substantial demand. Together they constitute something closer to a humanistic education than a technical skill.

Consider hallucination, the most discussed failure mode. A user with operational literacy knows to check facts. A user with comprehensional literacy understands that hallucination is not a bug to be patched but a structural feature of probabilistic text generation — that the model has no representation of truth, only of plausibility. China's first AI hallucination lawsuit, in which a student sued over fabricated university-consultation guidance that altered his Gaokao-related choices, became a public lesson precisely because most users had been operating on the first kind of literacy and discovered, painfully, that they needed the second [4]. The medical literature is now documenting what physicians are calling "never-skilling" — the phenomenon of practitioners who relied on AI scaffolding through training and never developed the underlying judgment to detect when the scaffolding is wrong [3]. Operational literacy without comprehensional literacy is a trap, and the trap closes most tightly on those who were taught only the operational layer.

[4] China's First AI Hallucination Lawsuit | Gaokao Error

[3] AI-induced never-skilling in medical education

Consider also the deepfake question. The Knight Columbia team's analysis of 78 election deepfakes found that the political-misinformation panic has been substantially miscast: the dangerous deepfakes are not the obvious ones that fool nobody but the mundane ones that erode the evidentiary baseline of public life [28]. Google's Veo 3 has made cinematic-quality fabrications of riots, conflict, and protest cheap and ambient [9]. The literacy task this poses is not "spot the fake" — by the time the spotting is needed, the epistemic damage is done — but "understand what it means to live in an environment where the cost of fabricating evidence has collapsed." UNESCO's analysis of deepfakes and the "crisis of knowledge" treats this as a question about the social infrastructure of credibility rather than about individual detection skills [17]. The scoping review of generative AI and misinformation in *AI & Society* arrives at a parallel conclusion: detection-focused literacy is necessary and insufficient, and the deeper task is rebuilding what the authors call "epistemic resilience" [7].

[28] We Looked at 78 Election Deepfakes. Political Misinformation Is Not an

[9] Google's Veo 3 Can Make Deepfakes of Conflict, Riots, More - TIME

[17] Les deepfakes et la crise du savoir - UNESCO

[7] Generative AI and misinformation: a scoping review of the role of

Epistemic resilience is not a personal skill. It is a property of a population that has access to shared evidentiary norms, working insti-

tutions of verification, and the comprehensional literacy to use them. None of those preconditions can be supplied by a six-hour vendor module on responsible prompting.

The Surveillance Substrate Nobody Teaches

There is a layer of AI literacy that almost no curriculum currently touches: the recognition that AI systems are, among other things, surveillance instruments, and that interacting with them is an act with consequences for one's data, one's exposure, and one's future treatment. American school districts now deploy AI monitoring tools — Gaggle, GoGuardian, Securly — that watch student activity on school-issued Chromebooks in ways students and parents rarely understand, flagging behavior and surfacing it to administrators, sometimes to law enforcement [13]. MIT Technology Review's reporting on bulk-data LLM surveillance describes how the same model architectures being marketed as productivity tools are being deployed by federal agencies to mine communications at a scale that was technically infeasible three years ago [14].

[13] How AI monitors school Chromebooks and what it means for privacy

[14] How LLMs could supercharge mass surveillance in the US

A literacy that does not include the surveillance layer is not literacy; it is the trained obliviousness that allows the surveillance to scale. Yet a survey of the most popular K-12 and corporate AI-literacy curricula would find this layer almost entirely absent. The curricula teach how to use the tools. They do not teach what the tools are doing while being used. There is a reason. Curricula taught by vendors, or by institutions whose vendor relationships are load-bearing, do not lightly include modules on how those vendors monetize the user.

The accessibility frame illustrates the same dynamic from a different angle. Tech Policy Press recently observed that the web is increasingly being optimized for AI agents rather than for human users, with structured-data accommodations made for crawlers and chatbots that have no equivalent investment for screen readers or low-bandwidth users [26]. The developer-facing version of the same story celebrates this as accessibility becoming "business-critical" because agents need clean markup [12]. The two framings describe the same phenomenon — the web is becoming machine-legible — and the only thing distinguishing them is whether you understand who the legibility serves. That understanding is what literacy is supposed to be for.

[26] The Web Is Being Made Accessible for AI, Not People

[12] How AI Agents Are Making Accessibility a Business-Critical Development Priority

The Geography of Who Counts

Whose literacy counts is decided largely by who is in the room when "literacy" is being defined. The OECD-funded Renaissance Numérique report from October 2025 was unusually candid about this, noting that French AI-literacy initiatives have been overwhelmingly shaped by industrial federations and grandes écoles, with civil-society participation reduced to consultation rather than co-design [6]. The Organization of Ibero-American States' mapping of AI's arrival in Latin American education is even blunter: regional capacity to define what AI literacy should mean has been outpaced by the speed of platform deployment, leaving ministries reacting to vendor-set agendas rather than setting their own [16].

[6] Déployer une littératie en IA pour une

[16] La llegada de la IA a la educación en América Latina: en construcción

This is the political-economic dimension that the operational definition of literacy is engineered not to see. When literacy is defined as "can you use the tool," the question of who designed the tool, in what language, with what assumptions about its users, becomes invisible by construction. The Generative AI Practices, Literacy, and Divides paper on arXiv documents how usage patterns vary across linguistic, geographic, and class lines in ways that map almost perfectly onto pre-existing structures of digital exclusion [8]. AI literacy is being scaled into a world where the people most in need of critical understanding of these systems are the ones with the least institutional voice in defining what understanding means.

[8] Generative AI Practices, Literacy, and Divides

The Spanish-language debate over whether schools should ban AI altogether captures the bind cleanly. The position is not Luddite; it is a recognition that absent the resources to teach critical literacy, the choice is between uncritical adoption and refusal, and refusal is the more defensible posture [29]. The Forbes call to block agentic AI browsers in schools arrives at a structurally similar conclusion from a security frame [5]. Both responses concede the same underlying problem: in the absence of comprehensional and civic literacy, operational fluency is a liability. The honest answer is to build the literacy. The realistic answer, given current investment, is to wall off the tools. Neither answer is what vendors are selling.

[29] ¿Deberían las escuelas prohibir la IA? La pregunta que no podemos

[5] Colleges And Schools Must Block And Ban Agentic AI Browsers ... - Forbes

The Democratic Stakes

The promise that AI could enrich democratic deliberation, rather than corrode it, depends entirely on which definition of literacy wins. Carnegie's recent work on AI-enabled deliberative democracy is openly conditional: the gains are real if the participating publics have the

capacity to evaluate AI contributions to deliberation, and illusory if they do not [24]. What this means in practice is that AI literacy is not a side issue in democratic life; it is increasingly the gate condition for democratic life. A polity in which most citizens cannot tell which of their inputs are being shaped by AI, which of their representatives' decisions are being scaffolded by AI, and which of the evidentiary objects in their public sphere were generated by AI is a polity that has lost the capacity for informed consent over a substantial fraction of its political process.

[24] Realizing the Potential Gains of AI-Enabled Deliberative

The deepfake-regulation map across twenty states demonstrates how slowly legal infrastructure adapts to this shift [11]. Regulatory patches address visible failures after the fact; literacy must operate before the fact, in the cognitive habits of citizens encountering ambiguous media. Journalists' guides to detecting AI-generated content, like the recent one from Verificado, do useful work, but their reach is limited to the professional verifier class [10]. The general public is asked to consume the verifiers' outputs on trust, which works only as long as trust in verifiers holds, which is exactly the resource AI-mediated misinformation depletes fastest.

[11] How 20 States Are Now Regulating Deepfakes—and What It Means for Elections

[10] Guía para periodistas sobre cómo detectar contenido generado por IA

Cross-domain, the same pattern recurs. The literacy gap is the gap that determines whether the equity claims made for AI in other domains — accessibility, healthcare, education, journalism — translate into actual equity. A special-education teacher using AI to draft IEPs operates on borrowed expertise: the tool is doing inference she may not have the time or training to audit, and the consequences of an unaudited inference accrue to a child [19]. A parent who sues after their child is disciplined for AI use confronts an administrative apparatus that has decided, on the basis of AI-generated suspicion, that their child has cheated [20]. In each case, the people most exposed to AI's consequences are the people least equipped, by current literacy definitions, to contest those consequences.

[19] Overworked and understaffed: Special ed teachers turn to AI for help

[20] Parents Sue After School Disciplined Student for AI ... - Education Week

What Would Count as Adequate

It is worth saying plainly what an adequate AI literacy would look like, because the field's failure to name it is part of how the inadequate version becomes the default. Adequate literacy would mean: a working understanding that these are probabilistic systems with no access to truth; a grasp of the labor and resource extraction underlying the models; the ability to recognize when one is being addressed by, monitored by, or evaluated by an AI system; a vocabulary for the surveillance and economic relations that one enters by using such sys-

tems; the capacity to evaluate vendor claims; and the civic standing to participate in decisions about deployment. These are not exotic demands. They are the AI-era analogue of the literacies that distinguished a citizen from a subject in the print era.

The MIT Press essential-knowledge volume on AI ethics frames the question in almost these terms, asking who will be empowered by AI and who will be excluded, and treating the answer as a matter of how a society chooses to define educational adequacy [27]. The book argues for an integrated Bildung in which doing AI and reasoning about AI's place in human life are not separate curricula. That is the right ambition. It is also a long way from what is actually being delivered in the workforce-development centers, the corporate-training portals, and the school districts where most actual literacy is being formed.

[27] AI Ethics - The MIT Press Essential Knowledge series

The French report's conclusion is worth taking seriously: literacy at the scale democracy requires cannot be built by vendor curricula or by isolated school initiatives. It requires public commitment, public funding, and public standards that are not set by the actors with the strongest interest in operational compliance [6]. UNESCO's work on critical thinking and AI gestures toward what such standards might look like, but its uptake remains spotty and its funding modest compared to the vendor-supplied alternatives [27]. The asymmetry of investment is the asymmetry of outcomes.

[6] Déployer une littératie en IA pour une

[27] UNESCO Think Critically Click Wisely

The Oracle's Mirror

Return to the figure of the oracle. Oracles in classical antiquity were not consulted by anyone who walked up; access was mediated by priesthoods, languages, and rituals that the consulting citizen had to learn or hire someone to perform. The democratic claim of the AI era was that the oracle would be available to all, that natural language would dissolve the priesthood. What has actually happened is more interesting and more troubling. The interface has become available to all, while the literacy required to use it without being used by it has been concentrated in a thin stratum, and that stratum is being trained — by the very institutions whose models it commands — to call its operational fluency "literacy" and to leave the comprehensional and civic layers for someone else to worry about.

The honest thing to say about AI literacy in late 2026 is that we have not yet decided, as a society, what the word means. The decision is being made for us, fragment by fragment, by the actors with the most coherent interest in a particular answer. A curriculum that teaches prompting without political economy, tools without surveil-

lance, use without critique, is not a neutral starting point that can be supplemented later. It is a finished product, and its finishing is what makes the supplementation unlikely. The reader who finishes this essay knowing that the word ”literacy” is currently doing three jobs at once, and that only one of them is being funded, has begun the literacy in question. That is meant as a small consolation, and also as a small warning. The oracle is in the room. The question of who learns to command it, and on whose terms, is the question of what kind of public the next decade is going to have.

References

1. 2024 The Dawn of the AI Era - Common Sense Media
2. AI and Democracy: Mapping the Intersections - Carnegie Endowment for
3. AI-induced never-skilling in medical education
4. China’s First AI Hallucination Lawsuit | Gaokao Error
5. Colleges And Schools Must Block And Ban Agentic AI Browsers ... - Forbes
6. Déployer une littératie en IA pour une
7. Generative AI and misinformation: a scoping review of the role of
8. Generative AI Practices, Literacy, and Divides
9. Google’s Veo 3 Can Make Deepfakes of Conflict, Riots, More - TIME
10. Guía para periodistas sobre cómo detectar contenido generado por IA
11. How 20 States Are Now Regulating Deepfakes—and What It Means for Elections
12. How AI Agents Are Making Accessibility a Business-Critical Development Priority
13. How AI monitors school Chromebooks and what it means for privacy
14. How LLMs could supercharge mass surveillance in the US
15. IBM Brings Its Most Advanced AI-Powered Security Portfolio to Clients, and is Strengthened by Ongoing Project Glasswing Work

16. La llegada de la IA a la educación en América Latina: en construcción

17. Les deepfakes et la crise du savoir - UNESCO

18. New sources of inaccuracy? A conceptual framework for studying AI

19. Overworked and understaffed: Special ed teachers turn to AI for help

20. Parents Sue After School Disciplined Student for AI ... - Education Week

21. Personnaliser l'apprentissage pour les étudiants handicapés à l'aide de

22. Prompt engineering as a new 21st century skill

23. Prompt engineering in higher education: a systematic review to help

24. Realizing the Potential Gains of AI-Enabled Deliberative

25. The largest study of AI use by undergrads is in, revealing disparities in access — and in cheating

26. The Web Is Being Made Accessible for AI, Not People

27. UNESCO Think Critically Click Wisely

28. We Looked at 78 Election Deepfakes. Political Misinformation Is Not an

29. ¿Deberían las escuelas prohibir la IA? La pregunta que no podemos