# Fluency Without Understanding: The Empty Promise of Current AI Literacy

Weekly Analysis — https://ainews.social

A teenager in Ohio can prompt a chatbot to write a college admissions essay in the cadence of Joan Didion, debug a Python script with three follow-ups, and generate a passable deepfake of a classmate before homeroom. By every operational metric the major frameworks use, she is AI-literate. She is also, by any meaningful definition of the word, not literate at all — she cannot tell you when the model is hallucinating, cannot identify the economic interests served by its default outputs, cannot say what the system does to her own capacity for sustained thought, and has never been asked to consider whether any given task should have been done with the tool at all. This is not her failure. It is the failure of a curricular consensus that has, with remarkable speed, mistaken comfort for comprehension.

The vendor frameworks now setting the terms of public understanding are quite explicit about what they want literacy to mean. Microsoft's [18] module organizes its competencies around productive use of Copilot, Azure services, and the company's reference architectures; its companion [6] frames the entire educational journey as one of "building practical skills" inside a vendor stack. There is nothing wrong with knowing how a tool works. There is something deeply wrong with allowing a tool's manufacturer to write the definition of the cognitive capacity citizens need to live alongside it.

[18] Introduction to AI Literacy - Training | Microsoft Learn

[6] AI learning hub - Start your AI learning journey, and build practical ...

What follows is a map of how "AI literacy" came to mean what it now means in most policy documents, training programs, and corporate learning paths — and what that map omits. The omissions are not accidental. They are the predictable result of letting the firms that profit from frictionless adoption define the meaning of competent use.

## *The Vendor's Definition and Its Discontents*

If you collect the AI literacy frameworks now circulating among employers, ministries of education, and large platform companies and lay them flat on a table, you find a striking convergence. They almost all distinguish four bands of competence: understanding what AI is, using AI tools effectively, evaluating AI outputs, and applying some

ethical considerations to AI use. Microsoft's own training path follows precisely this scaffold, moving from "what is generative AI" through "responsible AI principles" to hands-on prompting [18]. The European observatory data on student usage suggests this scaffold is succeeding on its own terms: French students surveyed in early 2026 report that they have integrated generative AI into nearly every stage of academic work, from initial research to final revision [23].

[18] Introduction to AI Literacy - Training | Microsoft Learn

[23] Observatoire des usages de l'intelligence artificielle par les étudiants

So what is the problem? The problem is that "effective use" and "evaluating outputs" are not the same skills as "knowing when not to use the tool" or "noticing what the tool is doing to your thinking." The vendor framework treats AI as a productivity layer to be optimized. It does not contemplate the user who, having reflected, declines. Pew's February 2026 survey of American teens found that roughly a quarter of those who use ChatGPT for schoolwork report doing so for tasks they acknowledge they could complete themselves; the question of whether they should have is not part of the literacy they have been taught [15].

[15] How Teens Use and View AI

The UNESCO frameworks try to push against this current. [11] is significantly more sophisticated than the vendor curricula, locating AI literacy within a longer tradition of media and information literacy and insisting that users develop competencies in recognizing AI, understanding intelligence as a contested concept, and analyzing algorithmic bias. The UNESCO [3] explicitly warns that generative outputs are "stochastic" and "potentially less trustworthy, especially for the teaching of factual and conceptual knowledge." This is closer to what literacy should mean. But it remains a minority report. The frameworks actually shaping how hundreds of millions of people are being taught to interact with these systems come from the firms selling them.

[11] Deepfakes and the crisis of knowing - UNESCO

[3] IA y accesibilidad: ¿renunciando al compromiso? - UNESCO

## *The Trick of Fluency*

There is a particular kind of confidence that AI tools produce in their users, and it is the central pedagogical problem of our moment. The output is articulate. The interface is friendly. The response arrives in seconds. Every textural cue tells the user that something competent has happened. Whether anything competent has actually happened is a separate question, and one most users are not equipped to ask.

The deepfake epidemic in secondary schools makes this concrete. Reporting from WIRED on the global spread of "nudifying" apps describes a generation of students who have technical fluency with image-generation tools but no developed framework for understanding

what those tools are doing to consent, evidence, or the social fabric of a school [25]. The Guardian's interactive investigation traced cases in which victims discovered fabricated images of themselves circulating among classmates who, when interviewed, often did not understand they had done anything that crossed a line — the tool made it easy, the output looked real, and no one had ever asked them to think about the gap between what was technically possible and what was ethically defensible [27]. The perpetrators were, in the vendor sense, AI-literate. They knew how to use the tool.

[25] The Deepfake Nudes Crisis in Schools Is Much Worse Than You Thought - WIRED

[27] The rise of deepfake pornography in schools: 'One girl was so horrified ...

This is what fluency without understanding produces. McLuhan, writing decades before any of this was possible, insisted that any medium reshapes what its users perceive as figure and ground, foregrounding some capacities and quietly atrophying others [22]. The generation now coming of age with chatbots as default cognitive partners is being trained to foreground output and background the question of how that output was produced, by whom, for whose benefit, and at what cost to their own thinking. The frameworks teaching them to use these tools rarely surface the ground at all.

[22] Understanding Media

The same trick operates in adult professional contexts. Tech Policy Press has documented how "human in the loop" arrangements — long sold as the ethical guardrail that makes AI deployment safe — frequently collapse into rubber-stamping, because the human reviewers face exactly the fluency problem the systems were designed to exploit [5]. When the machine's output looks confident and the queue is long, the human's role becomes performative. Literacy in this context would mean knowing when to refuse the assignment of pseudo-oversight. No vendor framework teaches that.

[5] AI Efficiency Can Undermine Accountability Even With ...

## *What the Frameworks Cannot See*

The deepest failure of current literacy curricula is not what they include but what they cannot conceive of including. Three blind spots organize the rest.

The first is the metacognitive blind spot — the inability to teach users to notice what the tool is doing to their own cognition. A growing body of research on student AI use, synthesized in a recent scoping review published in *AI & Society*, finds that heavy generative AI use correlates with measurable declines in the kinds of effortful reasoning that produce durable understanding, even when surface task performance improves [13]. Users get the answer; they lose the path to having found it themselves. A literacy framework worth the name would treat this as the central pedagogical concern. Most do not mention it.

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

The second is the political-economic blind spot. The frameworks treat AI as a neutral capability rather than as the product of a specific industry with specific interests in a specific configuration of public credulity. Shoshana Zuboff's analysis of how the data extraction economy produces "behavioral surplus" by treating user activity as raw material remains the indispensable account of what is actually happening when one of us types into a free chatbot [22]. A literacy framework that does not equip users to see this — to recognize that the seamlessness of the experience is itself a designed feature serving a business model — is a framework that produces compliant users rather than capable citizens. Kate Crawford's mapping of the material and labor substrates of AI systems makes the same point from a different angle: the cloud is not a cloud, the intelligence is not autonomous, and the costs are not where the marketing places them [22].

[22] The Age of Surveillance Capitalism

[22] The Atlas of AI

The third is the epistemic-authority blind spot. UNESCO's own analysis of the deepfake crisis describes how generative systems have produced what it calls a "crisis of knowing" — a generalized erosion of the public's capacity to determine what is real, what is fabricated, and who to trust [11]. The Reuters Institute's investigation of disinformation targeting Latino voters in the United States documented exactly this dynamic playing out in a specific democratic context, with AI-generated content amplifying conspiracy theories in Spanish-language information ecosystems where fact-checking infrastructure is thinner [16]. A literacy framework focused on "evaluating outputs" cannot get at this, because the unit of analysis is wrong — the problem is not whether any individual output is accurate but what the saturation of synthetic content does to the social process of knowing.

[11] Deepfakes and the crisis of knowing - UNESCO

[16] IA, mentiras y conspiranoia: así sufren la desinformación los votantes ...

## *The Asymmetry of Harm*

When literacy fails, it does not fail uniformly. The populations most exposed to AI-mediated harm are not the ones whose schools have invested in critical AI curricula or whose employers have built reflective practice into their training paths. They are the populations whose interactions with these systems are mediated by underfunded institutions, by chatbots whose deployments are calibrated to extract attention rather than build capacity, and by content streams optimized to inflame.

The Bureau of Investigative Journalism's reporting on a Pakistani entrepreneur monetizing Islamophobic AI slop on Western platforms describes a perfect closed loop: a content producer with no particular ideological stake generates synthetic outrage content because the plat-

forms reward it, the platforms reward it because users engage with it, the users engage with it because they cannot tell it from real reporting, and the targets of the content — Muslim communities — bear the social and political costs of an information environment they did not design [26]. At every node of this loop, the operative literacy is the wrong one.

[26] The devout Muslim making a living from Islamophobic AI slop

Ruha Benjamin's account of how technologies inherit and amplify the social structures of the societies that build them — what she names the New Jim Code — is the necessary frame here [22]. The deployment of AI is not neutral with respect to who bears its risks. The Markup's investigation of AI writing detectors found that the tools systematically misclassified work by non-native English speakers as machine-generated, producing a wave of false cheating accusations that fell hardest on international students [4]. Bloomberg's coverage of detector reliability reached the same conclusion from the institutional side: the tools do not work as advertised, but they are being used as if they do [12]. A peer-reviewed argument against generative AI detection in assessment makes the structural case: the entire enterprise rests on a false confidence in the discriminative power of the tools and produces disparate impact by design [10].

[22] Race After Technology

[4] AI Detection Tools Falsely Accuse International Students of Cheating

[12] Do AI Detectors Work? Students Face False Cheating Accusations - Bloomberg

[10] Contra generative AI detection in higher education assessments

Virginia Eubanks's documentation of how automated decision systems concentrate their failures on the poor — a population with the least institutional capacity to contest the systems' verdicts — describes the broader pattern of which the detector scandal is one instance [22]. The literacy gap is not just a knowledge gap. It is the cognitive analog of a power gap, and treating it as a curricular problem to be solved with better prompt engineering modules is a category error.

[22] Automating Inequality

## *The Youngest Users and the Oldest Failure*

The asymmetry sharpens further when the user is a child. Brookings, in a recent guide to the technologies already shaping young children's lives, points out that "Generation AI" begins long before any classroom curriculum reaches them — voice assistants, recommender systems, and adaptive learning platforms have been training their cognition from before they could read [14]. Whatever literacy we eventually offer them will be retrofit literacy, attempting to make conscious what has already been internalized as ground.

[14] Generation AI starts early: A guide to technologies already shaping ...

A taxonomy of generative AI risks to youth published at USENIX SOUPS 2025 catalogues the specific harms — from emotional dependence on chatbot companions to manipulation by recommender

feedback loops to the casual normalization of synthetic intimate imagery — and notes how few of these risks figure in mainstream literacy curricula at all [29]. Common Sense Media's 2025 report on teen trust in AI found that adolescents themselves report low confidence in the information chatbots produce — and continue to use them anyway, because the alternatives are friction and the chatbot is right there [1]. Distrust without behavioral change is the signature of a population that has been given the wrong tools to act on its instincts.

[29] Youth-Centered GAI Risks (YAIR): A Taxonomy of Generative AI Risks from ...

[1] 2025 Teens, Trust, and Technology in the Age of AI - Common Sense Media

The harms here are not hypothetical. Ipsos's European survey on conversational AI and youth mental health documented a pattern of adolescent reliance on chatbots for emotional support in ways that displace, rather than complement, human connection [20]. A Kentucky lawsuit recently described in Bloomberg Law offers a litigation blueprint for states attempting to hold AI chatbot companies accountable for harms to minors, an avenue that has emerged precisely because consumer-facing literacy interventions have so manifestly failed to mitigate the underlying dynamics [19]. CNN's coverage of an emerging "independent crash testing" lab for youth-facing AI products suggests the institutional vacuum is now being filled by external auditors, not by the literacy frameworks that were supposed to equip users to evaluate the systems themselves [9]. When the response to a literacy failure is a third-party crash-testing regime, the literacy frameworks have already lost the argument.

[20] L'IA conversationnelle et la santé mentale des jeunes en ...

[19] Kentucky Lawsuit Offers Blueprint for States to Sue AI Chatbots

[9] Child safety lab launching 'independent crash testing' for AI ...

A study from researchers at Illinois examining teen use of generative AI captured the texture of the gap directly: teens know enough to use the tools, know enough to be uneasy about some of what they encounter, and lack the conceptual scaffolding to translate that unease into informed action [17]. This is fluency without understanding in its purest form. The vendor frameworks produced exactly the outcome they were designed to produce.

[17] Illinois researchers examine teens' use of generative AI, safety ...

## *Civic Literacy and the Architecture of Consent*

The argument has so far focused on individual cognition, but the stakes are civic. A society of fluent users who cannot resist the cognitive demands of the tools they use is a society whose democratic deliberation has been outsourced to the platforms that mediate it.

The Springer Nature audit of TikTok's algorithm during a recent national election cycle found that the platform's recommender systems produced systematically skewed political content exposure in ways the platform itself had not disclosed and most users could not detect [8]. The literacy implication is not that users should learn to evaluate indi-

[8] Auditing TikTok's Algorithm During the Most Consequential ...

vidual posts — that framing concedes the wrong terrain. The literacy implication is that users need a working theory of how recommender architectures shape the universe of inputs from which any individual evaluation must proceed. Eli Pariser's account of how personalization produces epistemically isolating bubbles is now nearly fifteen years old; nothing in the mainstream literacy frameworks has caught up to its analytic level [22].

[22] The Filter Bubble

Noam Chomsky and Edward Herman's filter model of how mass media manufactures consent through ownership, advertising, sourcing, flak, and ideological framing was developed for an information environment that, in retrospect, looks artisanal compared to the one AI now mediates [22]. The filters have not disappeared; they have been automated, personalized, and rendered illegible. A literacy framework adequate to the present moment would treat understanding how the filters now operate as a basic civic competency. The vendor frameworks treat it as out of scope.

[22] Manufacturing Consent

Meredith Broussard's argument that the technology industry's "technochauvinism" — its insistence that technical solutions are always preferable to social ones — has produced a generation of systems whose failures are systematically misattributed to user error supplies the diagnostic frame [22]. When the literacy curriculum tells you that the appropriate response to a flawed AI output is better prompting, the curriculum is participating in the technochauvinism it should be helping you see through.

[22] Artificial Unintelligence

A systematic review on responsible AI in education published in *Nature*'s humanities and social sciences journal earlier this year reaches a similar conclusion from inside the educational-technology literature: most of what is currently sold as responsible AI integration consists of operational controls and disclosure mechanisms that do not meaningfully transfer epistemic agency back to the user [28]. The review's authors argue, with appropriate caution, that the field has not yet developed a vocabulary for what genuine literacy in this context would require. They are right. The vocabulary does not exist because the frameworks that have captured the field had no interest in developing it.

[28] Towards responsible artificial intelligence in education: a systematic ...

## *What Literacy Would Have to Become*

If the vendor definition is the wrong one and the existing critical frameworks have not yet been operationalized, what would adequate AI literacy actually require? The answer is not a longer competency checklist. It is a different conception of what literacy is for.

The first move is to reintroduce the unplugged. Any literacy curriculum that does not include sustained periods of working without the tool is teaching dependence, not capacity. This is not nostalgia. It is the basic insight that one cannot evaluate the cognitive effects of an instrument without spending time outside its influence — and one cannot recognize what the tool is doing to one's thinking without having an unaided baseline to compare it against. The metaliteracy tradition has been arguing for variants of this position for over a decade, framing it as the ability to consciously toggle between participatory and reflective modes of engagement with information environments [22].

[22] MetaLiteracy

The second move is to teach refusal as a literacy. Knowing when not to use a tool is a higher-order skill than knowing how to use it. A worker who can articulate, in front of their manager, why a given task should not be automated has more functional literacy than a colleague who can prompt their way through it faster. A student who declines to outsource an argument they wanted to develop themselves has more literacy than one who produces a polished essay they cannot defend. The educavox recension of disinformation-education initiatives across French-speaking Europe and Quebec captures the texture of curricula that take this seriously [30], as does the empirical work documented by Quebec researchers on empowering young people specifically to resist synthetic disinformation rather than merely identify it [21].

[30] Éduquer contre la désinformation amplifiée par l'IA et l'hypertrucage ...

[21] Les « deepfakes » : Comment donner aux jeunes les moyens de lutter ...

The third move is to teach the political economy of the tools as a first-order component of competence, not as an "ethics module" tacked on at the end. The user who understands that the chatbot is a commercial product whose default behaviors reflect business decisions has a different relationship to its outputs than the user who experiences it as a neutral oracle. Forbes's recent argument that AI's accessibility commitments will be tested by exactly the moments when accessibility cuts against engagement metrics points toward the kind of analysis users themselves would need to do to be genuinely literate [7]. The legal-regulatory literacy now becoming essential — captured in analyses like the Spanish-language treatment of how GDPR and the AI Act interact when students use AI in thesis preparation [24] — is another component the vendor curricula systematically omit, because the firms have no incentive to surface the constraints their products operate under.

[7] Artificial Intelligence Has One Chance To Get Accessibility Right

[24] Privacidad de Datos en Tesis con IA: Marco RGPD y AI Act 2026

The fourth move is to make metacognitive self-monitoring an explicit subject of practice. Users need vocabulary and exercises for noticing when the tool is doing their thinking for them, when their judgment is being smoothed into compliance, when the convenience of the system has begun to deform their sense of what counts as a

question worth asking. This is not mysticism. It is the same kind of disciplined attention that any serious craft tradition has always required of its practitioners. The MIT Press *AI Ethics* volume frames this as a question of what kind of *Bildung* — what kind of formative education — would be adequate to a technologically saturated society [22]. The honest answer is one we have barely begun to develop, and one the vendors who have written most of the existing curricula have no interest in producing.

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

## *Whose Literacy Counts*

The question that haunts this entire field is who gets to decide what literacy means. Microsoft's training paths reach hundreds of millions of working adults; UNESCO's frameworks reach ministries of education whose budgets cannot match the vendors' marketing spend; the critical scholars who have done the most rigorous work reach a few thousand readers per book. The framework that wins is not the framework with the best argument. It is the framework with the deepest distribution.

This is the asymmetry the reader needs to see clearly. When a corporation defines the competencies citizens need to deal with the corporation's products, the resulting definition will be one the corporation can supply training for. When a public-interest body defines those competencies, the definition expands to include capacities the corporation has no incentive to teach. The current moment is one in which the corporate definition has decisively captured the policy conversation in most jurisdictions, and the critical alternative survives in scattered academic and regulatory pockets. Chomsky's late-career insistence that the deepest victories of concentrated power are the ones that limit what counts as a serious question is the relevant frame [22]. The serious question right now is not "how do we teach people to prompt better." The serious question is "what kind of cognitive sovereignty do citizens need to retain, and what curricula would protect it."

[22] Necessary Illusions

A statistics report on AI cheating in schools recently estimated that the gap between students who have been taught to use AI tools and those who have been taught to think critically about them now correlates strongly with measurable differences in academic integrity outcomes, with the former group reporting higher rates of unreflective dependence and the latter showing more discriminating use [2]. The implication runs in the obvious direction. Whatever we are calling AI literacy now is producing the outcomes the critical traditions warned

[2] AI Cheating in Schools: 2026 Global Trends & Bias Risks

it would produce. The frameworks that diagnosed the problem in advance are the ones not being adopted.

The empty promise of current AI literacy is the promise that competence with a tool is the same as understanding what the tool is doing to the people who use it, to the institutions that deploy it, and to the social fabric that has to absorb its outputs. It is not the same. It has never been the same. The cost of pretending otherwise will be paid by populations who were told they had been made ready, and who discover, too late, that the readiness they were sold was the readiness their tool's manufacturers needed them to have. A literacy worthy of its name would have taught them to see that coming. We still have time to build one. We will not do it by accepting the definitions on offer.

## *References*


1. 2025 Teens, Trust, and Technology in the Age of AI - Common Sense Media
2. AI Cheating in Schools: 2026 Global Trends & Bias Risks
3. IA y accesibilidad: ¿renunciando al compromiso? - UNESCO
4. AI Detection Tools Falsely Accuse International Students of Cheating
5. AI Efficiency Can Undermine Accountability Even With ...
6. AI learning hub - Start your AI learning journey, and build practical ...
7. Artificial Intelligence Has One Chance To Get Accessibility Right
8. Auditing TikTok's Algorithm During the Most Consequential ...
9. Child safety lab launching 'independent crash testing' for AI ...
10. Contra generative AI detection in higher education assessments
11. Deepfakes and the crisis of knowing - UNESCO
12. Do AI Detectors Work? Students Face False Cheating Accusations - Bloomberg
13. GenAI and misinformation in education: a systematic scoping ... - Springer
14. Generation AI starts early: A guide to technologies already shaping ...

15. How Teens Use and View AI

16. IA, mentiras y conspiranoia: así sufren la desinformación los votantes ...

17. Illinois researchers examine teens' use of generative AI, safety ...

18. Introduction to AI Literacy - Training | Microsoft Learn

19. Kentucky Lawsuit Offers Blueprint for States to Sue AI Chatbots

20. L'IA conversationnelle et la santé mentale des jeunes en ...

21. Les « deepfakes » : Comment donner aux jeunes les moyens de lutter ...

22. Necessary Illusions

23. Observatoire des usages de l'intelligence artificielle par les étudiants

24. Privacidad de Datos en Tesis con IA: Marco RGPD y AI Act 2026

25. The Deepfake Nudes Crisis in Schools Is Much Worse Than You Thought - WIRED

26. The devout Muslim making a living from Islamophobic AI slop

27. The rise of deepfake pornography in schools: 'One girl was so horrified ...

28. Towards responsible artificial intelligence in education: a systematic ...

29. Youth-Centered GAI Risks (YAIR): A Taxonomy of Generative AI Risks from ...

30. Éduquer contre la désinformation amplifiée par l'IA et l'hypertrucage ...