

The Literacy of Fear: Scams, Deepfakes, and the Missing Collective Response

Weekly Analysis — <https://ainews.social>

The most frequent verb in this week's AI literacy coverage is not *learn*. It is *protect*. Read enough of it and a picture forms of the ideal AI-literate citizen: someone who hangs up on a cloned voice, second-guesses a video, and hovers a suspicious cursor before clicking. This person is vigilant, skeptical, and profoundly alone. They have been taught to fear the technology and to guard their bank balance against it, but almost nothing in the discourse invites them to ask who built the machine, who profits from its errors, or how a public might govern it together. The literacy on offer is a literacy of fear, and fear, as a curriculum, has a shape. It individualizes. It monetizes. It stops precisely where collective power would begin.

This is not an accident of tone. It follows from who is doing the teaching. When a bank publishes a guide to AI voice-cloning fraud, it frames the threat as theft — your money, your account, your family's savings — and the remedy as personal caution. The [3] genre teaches you to establish a "safe word" with relatives and to call back on a known number, which is sensible advice as far as it goes. What it never does is treat the same technology as a threat to your autonomy, your capacity to know what is true, or your standing as a participant in public life. The National Council on Aging's advice on [13] is pitched, reasonably, at protecting older adults from being fleeced — but the frame is consumer defense, not democratic understanding. The literate subject these texts imagine is a potential victim to be hardened, not a citizen to be empowered.

I want to map the muddle that results — the competing definitions of "AI literacy" jostling in this week's coverage, the skills-versus-critique tension underneath them, and above all what the dominant frame quietly excludes. Because the stakes are not abstract. If literacy means only personal safety and consumer vigilance, we will produce a public of suspicious shoppers, exquisitely trained to protect their wallets and utterly unequipped to shape the social role of the systems shaping them.

[3] AI Voice Cloning Scam: How Fraudsters Fake Your Family's Voice

[13] How to Spot Deepfake Scams and Protect Yourself From AI Fraud

The Literacy They Are Selling You

Start with the sheer volume of the fraud frame, because volume is itself an argument. This week's most legible literacy lesson is: AI can imitate people to steal from you. Spanish-language fact-checkers explain [5], walking readers through the mechanics of synthetic-voice cons. French security writers dissect the [7], the "CEO fraud" in which a cloned executive voice authorizes a fraudulent wire transfer. The problem is real, the harm is documented, and the advice is often good. But notice the consistent grammar of these pieces: the peril is financial, the target is you, and the response is a set of habits you adopt privately. Nowhere does the reader graduate from protecting an account to interrogating an industry.

This is what it looks like when literacy is authored by the parties whose interests it serves. Chomsky and Herman's account of how institutional filters shape the news is useful here not as a conspiracy but as a structural observation: the messages a society receives about power tend to be produced by, and to protect, that power. When banks and security vendors become the primary educators about AI, the syllabus predictably centers the risk they are equipped to manage — transactional fraud — and omits the risks they are not, such as the erosion of shared reality or the concentration of algorithmic control. [23] describes exactly this narrowing of the frame: not censorship, but a definition of the relevant problem so tight that certain questions never get asked. The bank does not have to suppress the autonomy question. It simply never raises it, and its silence becomes the shape of your literacy.

The consumer frame runs deeper than any single article, because it draws on a subjectivity the technology industry has spent two decades cultivating. Zuboff's argument that we have been reconstituted as data-generating consumers rather than as citizens is directly relevant to how "literacy" now gets defined; a research analysis of [24] the self traces how thoroughly the market frame has colonized our self-understanding. [23] names the mechanism: when the dominant institutions relate to you as a consumer, they teach you to relate to yourself the same way. So the "AI-literate" person emerges as a careful buyer of digital safety, comparison-shopping among password managers and verification apps, rather than a member of a public with a claim on how these systems are designed and deployed. The pedagogy of fear is, at bottom, a pedagogy of consumption.

And fear sells a particular product: the aftermarket detector. The reader trained to be afraid is then sold the tool that promises to allay the fear — the deepfake-spotting service, the identity-monitoring

[5] Cómo se usa la inteligencia artificial (IA) en 'deepfakes' para estafar

[7] Deepfake vocal + virement : la nouvelle fraude au président

[23] Manufacturing Consent

[24] Surveillance Capitalism, Consumer Subjectivity and

[23] The Age of Surveillance Capitalism

subscription. This is a closed loop, and it forecloses the more uncomfortable questions. Why is non-consensual synthetic imagery so abundant that a report can find [6]? That figure indicts an industry and a permissive deployment culture, not the vigilance of individual users. Yet the consumer frame converts a structural scandal into a personal hazard to be managed. You are taught to watch for the fake, not to ask why the machinery for producing it was released, unmarked and at scale, into the commons.

Forensics Without the Public

If the first flaw in the fear curriculum is that it individualizes, the second is that it withholds the very tools it implies you need. The dominant advice is to *be suspicious* — but suspicion without method is just anxiety. Tell people that any video might be fake and you have not made them literate; you have made them disoriented, which is a different and more manipulable condition. UNESCO’s essay on [8] names this precisely: the deepest damage is not any single forged clip but the general collapse of confidence in evidence, the “liar’s dividend” in which the powerful can dismiss authentic footage as fabricated. Generalized doubt is not critical thinking. It is the raw material for cynicism, and cynical publics are easy to govern.

What would forensic depth actually require? It exists, but it lives mostly outside the popular literacy discourse, in technical and institutional venues that ordinary readers never encounter. Researchers have built, for instance, a [18], capable of tracing a synthetic clip back to the generative model that produced it — provenance as a research problem, not a consumer reflex. Governments and standards bodies have advanced cryptographic provenance schemes; the NSA-affiliated guidance on [19] describes signing media at capture so its origins travel with it. The open-source intelligence community has its own reckoning, mapped in [14], about how generative systems corrode source verification and what disciplined verifiers do in response. This is real literacy — provenance, chain of custody, model attribution, cross-checking. And almost none of it reaches the reader of the bank’s scam guide.

The gap is instructive. The forensic knowledge is being developed as infrastructure, for institutions, while the public is handed vibes — “trust your gut,” “look for weird hands.” Broussard’s insistence that we demystify the machine rather than mythologize it applies directly here; [23] argues that treating computation as inscrutable magic disempowers people who could, with honest explanation, understand how

[6] Deepfake Crisis 2025: 96% Is Non-Consensual Porn

[8] Deepfakes and the crisis of knowing - UNESCO

[18] New tool identifies the sources of fake video | UCR News

[19] PDF Content Credentials: Strengthening Multimedia Integrity in the ...

[14] IA y OSINT: cómo la IA mina la verificación de fuentes

[23] Artificial Unintelligence

these systems actually work and fail. When the popular literacy stops at “be careful” and reserves the how-to-verify knowledge for defense agencies and research labs, it enacts exactly the mystification Broussard warns against. The public gets fear; the institutions get forensics.

There is a further asymmetry worth naming. Provenance systems like content credentials solve the authenticity problem from the *producer’s* side — they help trusted institutions prove their content is real. They do far less for the reader confronting an unlabeled clip in a feed, which is where most people actually meet the problem. European broadcasters wrestling with whether AI content should even be labeled — see the RTBF analysis of [9] — reveal how far we are from a public-facing standard. Labeling is a governance decision, made by platforms and regulators, that determines what the ordinary reader can even see. And that is the point the fear curriculum obscures: your ability to verify is not primarily a function of your personal alertness. It is a function of infrastructure that others control. A literacy that never says so leaves you responsible for a task the system has been designed to make impossible.

[9] Des contenus générés par IA
sources de plus en plus récurrentes de
...

People Before Algorithms, But How?

Here is where the discourse gestures at something better and then declines to develop it. The most promising counter-frame this week is collective rather than individual — the recognition that understanding AI is a shared civic project, not a private survival skill. You can see its outline in efforts to rebuild trust as a social good rather than a consumer purchase. A Queen’s University project asking [4] treats the problem as one of shared epistemic norms, not individual gullibility. Microsoft’s account of [1] — however much one should read a vendor’s self-description skeptically — at least names the possibility that a *community*, and not just a consumer, could hold power over how AI systems are trained and governed.

[4] Can This New Way to Fight
Misinformation Work?

[1] A new approach to AI data puts
communities in charge

The most robust version of this collective vision appears in the language of “people before algorithms” that has become a rallying cry in labor and civic circles. It is a genuinely different literacy claim: not *how do I avoid being scammed* but *how do we, together, decide what these systems are allowed to do to us*. The slogan is right. The problem is that it remains a slogan. This week’s coverage repeats the demand and never once examines the machinery that would fulfill it. When a labor organization insists that human judgment come before automated decision-making, what concrete algorithmic understanding does it bring to the negotiating table? How does a collective actually

learn to read a scheduling algorithm, audit a productivity-surveillance model, or contest a risk score? The coverage does not ask, and the silence is the story. We are told collective literacy is desirable without being shown that it can be built.

This is the missing pedagogy, and its absence has consequences that reach into governance itself. The French essay on [10] makes the sharpest version of the case: explicability — the capacity to understand and contest how an algorithmic system reaches its decisions — is not a technical nicety but a critical resource of the state and of citizenship, the very thing that allows a “social contract” to survive automation. Explicability is a *collective* competence. A single vigilant consumer cannot demand it; a union, a regulator, a civic coalition can. Yet the literacy discourse keeps returning to the lone individual and their private defenses, as though the social contract could be renegotiated one careful shopper at a time.

The Latin American media-education tradition has been more honest about what a collective, political literacy demands. The essay [12] argues for a “technopolitical” literacy — one that treats reading technology as inseparable from reading power, and that locates the learner within collective structures rather than as an isolated consumer of safety tips. This is the Freirean inheritance: literacy as the practice of naming and acting on the world together, not the private accumulation of defensive habits. The distance between “establish a safe word with your relatives” and “learn to collectively contest an algorithmic decision” is the distance between the literacy we are being sold and the literacy a democracy requires.

The Machine That Profits From Being Wrong

Nothing exposes the shallowness of the dominant literacy quite like the way it explains AI’s errors. When popular coverage bothers to explain why these systems produce falsehoods, it reaches for technical vocabulary that sounds like understanding while carefully avoiding the part that would implicate an industry. The French explainer [22] is representative and genuinely useful: it attributes model errors to training-data quality, to overfitting, to the statistical nature of prediction. All true. All incomplete in the same telling way. The explanation stops at the machinery and never reaches the market. It treats hallucination as a bug in the code rather than a predictable output of an economy that rewards shipping fast, cheap, and confident over shipping slow, expensive, and reliable.

This omission matters because it hands the reader a mistaken the-

[10] Effondrement ou reconquête du contrat social algorithmique

[12] Hacia una Alfabetización Crítica y Tecnopolítica en el Aula ...

[22] Pourquoi l’IA fait des erreurs et comment les reconnaître

ory of the problem. If hallucination is merely a data-quality issue, then the fix is technical and the responsibility is diffuse — better data, more training, and the problem recedes on its own. But if hallucination persists partly because reliability is costly and speed-to-market is lucrative, then the reader is looking at a *choice*, made by firms, under competitive pressure, that could be regulated or resisted. The French policy analysis [20] situates generative error within an information *economy*, where the incentives of production and distribution — not just the internals of the model — determine how much falsehood floods the ecosystem. That is the missing half of the explanation, and its absence is not neutral. To omit the profit motive is to protect it.

[20] PDF IA générative et mésinformation

The stakes become concrete when these unreliable systems are consulted for decisions that used to be the province of shared civic knowledge. When researchers tested chatbots as sources of electoral guidance, they found the [11] — a failure with obvious democratic consequences that no amount of individual scam-vigilance addresses. The problem is not that a fraudster is impersonating your grandson; it is that a trusted-seeming system, deployed for profit and optimized for fluency over accuracy, is quietly corroding the informational basis of self-government. A literacy that trains you to guard your bank account has nothing to say about this. It has defined the threat as theft and thereby rendered itself blind to the threat as epistemic capture.

[11] Election voting advice from AI chatbots 'inaccurate and unreliable'

The primer on AI ethics from the MIT Press Essential Knowledge series is blunt about the hidden human and economic scaffolding that the technical frame obscures — the miners, the ship workers, the underpaid data-labelers whose labor makes these systems possible, all "in the service of capital accumulation by very few people." [23] insists that "doing AI" honestly must include this political economy rather than bracketing it as someone else's concern. A literacy that explains hallucination through overfitting alone, while never mentioning the incentive structure or the labor beneath it, is not neutral technical instruction. It is the industry's preferred story about its own failures, repeated back to the public as education.

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

Whose Literacy Counts

Every definition of literacy encodes an assumption about who the learner is and what they are being prepared for. So it is worth asking, of this week's coverage, whose literacy is being built and whose is being neglected. The fear curriculum, for all its talk of protecting the vulnerable, has a curiously narrow imagination of the vulnerable — usually a solitary account-holder, often an older adult framed as a

soft target. It rarely imagines the learner as a worker facing an opaque management algorithm, a citizen facing an automated benefits determination, or a community facing a system trained on data it never consented to provide. These are the sites where AI's power is most consequential and where literacy would most need to be collective, and they are precisely the sites the dominant frame leaves dark.

Consider how differently literacy looks when the frame shifts from consumer safety to democratic participation. The Forbes essay [2] asks what it means to form citizens — not consumers — in an age of capable machines, which is the right question even when the answers remain thin. The French analysis of [17] probes what these systems implicitly teach about knowledge, authority, and truth — a formation that happens whether or not anyone designs it. These pieces point toward a literacy of judgment and participation rather than defense. But they remain minority reports against a torrent of scam guides, and the imbalance tells you which definition currently commands the field.

There is also a question of who gets counted as a full participant at all. UNESCO's work on media and information literacy has long insisted that literacy and equality are inseparable — that a person's ability to access, question, and act on information is conditioned by their social and economic standing. [23] makes the point through the lens of gender: women's unequal access to information technology "hinders their access and opportunity to express themselves and participate fully and meaningfully in social and economic development." The consumer frame is blind to this, because it presumes a generic user with money to protect and habits to adjust. It cannot see that the capacity for AI literacy is unequally distributed, and that a literacy worthy of the name would have to address that inequality as part of its content, not treat it as someone else's file.

The disinformation landscape sharpens the point about power still further, because the most sophisticated manipulation is not aimed at your bank account at all. It is aimed at your politics, and it is well funded. The Wall Street Journal's reporting on [16] describes influence operations of a scale and coordination that no individual verification habit could withstand. When a minister in India files a police complaint after a fabricated clip circulates — the case of [21] — the recourse available to the powerful is the recourse the ordinary citizen lacks: institutional, legal, collective. The asymmetry is the whole point. Coordinated, capitalized manipulation meets the individual's private skepticism, and the individual loses. UNESCO's broader treatment of [15] frames the response as necessarily systemic — governance, platform accountability, public infrastructure — precisely

[2] AI In Education: Forming Citizens When AI Can Do The Work

[17] Lo que la IA enseña (y no enseña) a los niños

[23] Think Critically Click Wisely

[16] Israel's \$50 Million Experiment to Change U.S. Public Opinion

[21] Piyush Goyal files police complaint over AI-generated deepfake video, warns of strict legal action

[15] Inteligencia artificial y desinformación - UNESCO

because the individual frame cannot bear the weight.

Benjamin’s analysis of how ostensibly neutral technologies encode and amplify existing hierarchies is the missing lens in most of this coverage. [23] argues that systems marketed as objective often reproduce discrimination under a veneer of technical impartiality — which means that “reading” AI critically requires reading for power, not just for authenticity. A literacy that teaches you to detect a forged voice but not to detect a biased risk score has trained you to spot the crude fraud while missing the systematic one. The scam is legible because it is a discrete crime with a discrete victim. The structural harm is illegible within the fear frame because it has no single villain and no safe word — only a design, a deployment, and a distribution of consequences that falls hardest on those the discourse never imagined as its learners.

[23] Race After Technology

The Literacy We Would Choose

Put the pieces together and the muddle resolves into a choice. Two definitions of AI literacy are competing for the public’s mind, and only one of them is being funded, published, and repeated at scale. The first — the literacy of fear — defines the learner as a consumer, the threat as theft, the remedy as private vigilance, and the horizon as personal safety. It is authored largely by the institutions whose interests it serves, it withholds the forensic depth it implies you need, it explains AI’s failures in terms that exonerate the market, and it leaves the structural harms — epistemic collapse, algorithmic control, encoded inequality — entirely off the page. It is not useless. Knowing how to hang up on a cloned voice is worth knowing. But as a definition of what it means to be literate in a world saturated by AI, it is a cage dressed as a shield.

The second definition — call it the literacy of participation — defines the learner as a citizen and a member of collectives, treats understanding as a shared resource rather than a private purchase, insists on explicability as a civic entitlement, and reads technology as an expression of power that publics have the standing to contest. Its fragments are visible this week in the technopolitical pedagogy of [12], in the argument that explicability is the [10] that keeps a social contract alive under automation, and in UNESCO’s insistence that literacy and equality rise or fall together. But these remain scattered, under-resourced, and — crucially — undertheorized in their collective dimension. Even the most promising counter-frame, the demand that people come before algorithms, has not been shown building the algorithmic understanding it would require to be more than a wish.

[12] Hacia una Alfabetización Crítica y Tecnopolítica en el Aula ...

[10] Effondrement ou reconquête du contrat social algorithmique

That is the missing collective response, and naming its absence is the beginning of demanding it. A public that can only protect its bank account is not a public capable of shaping AI's social role; it is a market segment being managed. The forensic knowledge exists — the provenance tools, the model-attribution research, the verification disciplines — and it is being built as infrastructure for institutions rather than as capacity for citizens. The political economy of error is well understood by those who study it, and systematically omitted from what the public is told. The collective vehicles for algorithmic literacy — unions, community data governance, civic coalitions — are being invoked as slogans without being examined as pedagogies. The gap between what we could teach and what we are teaching is not a knowledge problem. It is a power problem, and the fear curriculum is one of the ways that power keeps the public exactly literate enough to protect its money and exactly ignorant enough to leave the system unchallenged. The task is to insist on the other literacy — the one that treats you not as a target to be hardened but as a participant with a claim on the machine — and to notice, every time the discourse offers you a safe word instead, what it is quietly declining to teach.

References

1. A new approach to AI data puts communities in charge
2. AI In Education: Forming Citizens When AI Can Do The Work
3. AI Voice Cloning Scam: How Fraudsters Fake Your Family's Voice
4. Can This New Way to Fight Misinformation Work?
5. Cómo se usa la inteligencia artificial (IA) en 'deepfakes' para estafar
6. Deepfake Crisis 2025: 96% Is Non-Consensual Porn
7. Deepfake vocal + virement : la nouvelle fraude au président
8. Deepfakes and the crisis of knowing - UNESCO
9. Des contenus générés par IA sources de plus en plus récurrentes de ...
10. Effondrement ou reconquête du contrat social algorithmique
11. Election voting advice from AI chatbots 'inaccurate and unreliable'
12. Hacia una Alfabetización Crítica y Tecnopolítica en el Aula ...

13. How to Spot Deepfake Scams and Protect Yourself From AI Fraud
14. IA y OSINT: cómo la IA mina la verificación de fuentes
15. Inteligencia artificial y desinformación - UNESCO
16. Israel's \$50 Million Experiment to Change U.S. Public Opinion
17. Lo que la IA enseña (y no enseña) a los niños
18. New tool identifies the sources of fake video | UCR News
19. PDF Content Credentials: Strengthening Multimedia Integrity in the ...
20. PDF IA générative et mésinformation
21. Piyush Goyal files police complaint over AI-generated deepfake video, warns of strict legal action
22. Pourquoi l'IA fait des erreurs et comment les reconnaître
23. Race After Technology
24. Surveillance Capitalism, Consumer Subjectivity and