AI Literacy for Citizen Participation

Weekly Analysis — https://ainews.social

When UNESCO warns of a "crisis of knowing" in the age of artificial intelligence, they're not merely describing a technical problem. They're naming something deeper: the fracturing of our shared epistemic foundations. As [8] argues, we face not just the challenge of distinguishing real from synthetic, but the more fundamental question of how we know what we know. This crisis ripples through every domain of civic life, from classrooms to courthouses, from newsrooms to voting booths. Yet our response—the rush to develop "AI literacy"—reveals a conceptual battlefield where competing visions clash over what citizens need to know, who decides what counts as literacy, and whose interests these frameworks serve.

[8] Deepfakes and the crisis of knowing - UNESCO

The very term "AI literacy" has become a floating signifier, claimed by technologists pushing coding skills, ethicists demanding critical frameworks, corporations selling training programs, and governments drafting regulations. Each group maps the terrain differently, highlighting certain peaks while obscuring others. The result is less a coherent educational agenda than a series of overlapping and often contradictory projects, each claiming the mantle of preparing citizens for an AI-mediated future.

Understanding this contested terrain matters because literacy frameworks aren't neutral. They encode assumptions about power, agency, and participation. They determine whether we're training compliant users or critical citizens, whether we're focused on individual skills or collective capacities, whether we're addressing symptoms or structures. As we'll see, the battle over AI literacy is ultimately a battle over the kind of democracy we're building.

The Definitional Muddle: Three Competing Visions

The landscape of AI literacy frameworks reveals three distinct camps, each with its own logic and blind spots. The first camp champions functional literacy—the ability to use AI tools effectively. Here we find [1] proposing technical competencies around data handling, prompt crafting, and output evaluation. This approach treats AI as a tool to be mastered, emphasizing practical skills that enable productive use. Stanford's [22] expands this slightly, identifying four dimensions—

[1] AI & Data Competencies: Scaffolding holistic AI literacy in Higher Education

[22] Understanding AI Literacy | Teaching Commons

functional, ethical, rhetorical, and pedagogical—but still centers on enabling "effective and responsible use."

The functional camp's appeal is obvious: it promises immediate applicability. Students learn to write better prompts, professionals optimize workflows, teachers integrate AI assistants. The NSF's investment of $11 million in K-12 teacher training, documented in [NSF invests $11M to expand AI professional development for K-12 teachers nationwide], follows this logic—scale up practical training, democratize access to skills. Yet this approach risks reducing literacy to mere tool proficiency, creating what [20] calls "sophisticated users" rather than critical thinkers.

[20] Prompt Engineering for Education - Substack

The second camp pushes beyond functionality toward critical AI literacy. [10] exemplifies this approach with its comprehensive AILit framework, encompassing not just technical skills but "AI thinking," ethical reasoning, and societal impact assessment. This vision, echoed in [21], insists that true literacy requires understanding AI's role in reshaping power structures, labor markets, and democratic processes. Students don't just learn to use AI; they learn to question it.

[10] Empowering Learners for the Age of AI

[21] Towards an AI-Literate Future: A Systematic Literature Review - Springer

The critical literacy camp addresses what functional approaches miss: the political economy of AI, the reproduction of bias, the concentration of power in tech corporations. [11] demonstrates this expanded scope, examining how AI systems encode values and shape educational possibilities. Yet even critical frameworks can remain individualistic, focusing on personal ethical reflection rather than collective action.

[11] Ethics and governance of generative AI in education: a systematic...

The third camp—call it protective literacy—emerges from urgent concerns about AI-amplified harms, particularly to children. [17] typifies this approach, emphasizing defensive skills: recognizing deepfakes, understanding data harvesting, resisting manipulation. This camp's frameworks, evident in guides like [13], prioritize harm reduction over empowerment, teaching citizens to navigate threats rather than shape systems.

[17] ONU : Les risques croissants des contenus IA pour les enfants

[13] How can we keep children safe as AI reshapes the internet?

Skills Versus Understanding: The Fundamental Tension

The fault line between skills-based and understanding-based approaches to AI literacy runs deeper than pedagogical preference—it reflects fundamentally different theories of citizenship and agency. The skills camp, represented by initiatives like [14], operates on an implicit theory of adaptation. Citizens need skills to thrive in an AI-transformed world that is largely taken as given. The focus is pragmatic: how to use AI tools effectively in education, work, and daily

[14] Penn State launches AI literacy framework to advance University-wide AI literacy initiative

life.

This pragmatism has institutional momentum. [4] documents how quickly skills-based frameworks spread through educational systems, driven by employer demands and student anxieties about workforce readiness. The speed of adoption—[First-year SBU students to be required to take AI literacy] reports mandatory courses appearing across universities—suggests skills-based literacy fills a perceived urgent need.

[4] AI Literacy in K-12 and Higher Education in the Wake of Generative AI...

Yet the understanding-based camp, drawing on critical pedagogy traditions, questions what this urgency obscures. As [12] argues, focusing on skills without examining power structures risks creating "a new form of technical determinism." This French analysis proposes a tripartite governance model—institutional, teacher, student—that embeds critical questioning at every level. Similarly, Quebec's comprehensive guide [19] interweaves technical guidance with sustained ethical reflection.

[12] Gouvernance Algorithmique et Écosystèmes Éducatifs : Vers une Triade...

[19] L utilisation pédagogique, éthique et légale de l intelligence...

The tension becomes concrete in classroom practices. Should students first learn to use AI tools effectively, then develop critical perspectives? Or does starting with tool use naturalize AI's presence before critical faculties engage? [2] reveals this isn't merely theoretical—different sequences produce measurably different outcomes in student attitudes toward AI governance and regulation.

[2] AI adoption in schools global trends, benefits, and challenges

The most sophisticated frameworks attempt synthesis, but tensions persist. [1] proposes "holistic" integration, yet its learning outcomes still separate "technical competencies" from "ethical considerations," suggesting these remain distinct domains rather than interpenetrated aspects of literacy. The framework's very structure reveals the difficulty of true integration.

[1] AI & Data Competencies: Scaffolding holistic AI literacy in Higher...

The Epistemological Crisis: What Frameworks Miss About Knowing

Current AI literacy frameworks, for all their sophistication, largely miss what may be the deepest challenge: AI's disruption of how we establish truth and trust. [13] introduces the concept of "epistemic vertigo"—the disorienting loss of stable ground for determining what's real. This isn't simply about identifying deepfakes; it's about the erosion of the very frameworks we use to assess credibility.

[13] How AI swarms manipulate public opinion and how to fend them off

The crisis manifests in what researchers call the "liar's dividend"—the ability of bad actors to dismiss authentic evidence as potentially AI-generated. [23] documents how this dynamic operates in practice, finding that the mere possibility of AI manipulation undermines trust

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

in all media. Traditional media literacy, focused on source evaluation and fact-checking, proves inadequate when sources themselves become uncertain.

Most literacy frameworks acknowledge this challenge peripherally—mentioning deepfake detection or source verification—without grasping its depth. They offer technical solutions (detection tools, blockchain verification) to what is fundamentally an epistemological problem. As [14] reports from its public education efforts, citizens express less concern about identifying fake content than about losing any basis for shared truth.

[14] MSU Museum panel teaches about AI, politics and misinformation

The French framework [18] comes closest to addressing this crisis directly, proposing that AI literacy must include "epistemological reflection" on the nature of knowledge itself. Yet even this sophisticated approach struggles to operationalize such abstract concepts for practical education. How do you teach epistemology to eight-year-olds—or for that matter, to their teachers?

[18] Cadre pour l'utilisation pédagogique de l'intelligence artificielle...

The consequences ripple beyond individual confusion. [6] reveals how uncertainty about AI-generated content affects economic decision-making, while political scientists document democracy's dependence on shared factual baselines. When citizens can't agree on what's real, deliberation becomes impossible. The frameworks' failure to center this crisis represents their deepest limitation.

[6] Anthropic Economic Index report: Learning curves

Literacy as Power: Who Gets to Define Competence

The question of who defines AI literacy reveals the power dynamics current frameworks often obscure. [3] documents how definitions emerge primarily from three sources: technology companies, educational institutions, and government agencies. Each brings particular interests that shape what counts as literacy.

[3] AI Governance in Higher Education: Policy, Ethics & Risk

Technology companies, unsurprisingly, emphasize frameworks that position their products as essential. Google's [5] gamifies AI education in ways that naturalize corporate AI systems as the default environment children must navigate. The framework teaches critical thinking about AI-generated content while taking the existence and legitimacy of these systems as given. Microsoft and OpenAI's educational initiatives follow similar patterns—critical thinking within boundaries that don't question fundamental architectures of surveillance and control.

[5] AI Quests from Google teaches AI literacy to kids

Educational institutions, as [16] analyzes, face conflicting pressures. They must prepare students for an AI-mediated economy while maintaining academic integrity, promote innovation while protecting

[16] Making AI work for schools - Brookings

privacy, embrace efficiency while preserving human judgment. These tensions produce frameworks that often read as compromises—neither fully critical nor entirely accommodating, satisfying no one completely.

Government frameworks reveal different priorities. The comprehensive Spanish guide [18] emphasizes compliance, safety, and standardization. Quebec's framework, documented in [19], adds cultural preservation and linguistic sovereignty to the mix. These aren't neutral educational goals but political projects, using literacy frameworks to shape AI's integration according to state interests.

[18] Cadre pour l'utilisation pédagogique de l'intelligence artificielle...

[19] PDF L utilisation pédagogique, éthique et légale de l intelligence...

Notably absent from definitional power are students themselves, particularly those from marginalized communities. [Making Space: $500,000 Grant From GitLab Foundation To Expand AI Tools And Workforce Pathways For Disabled Talent] represents a rare exception, centering disabled perspectives in defining what AI literacy means for accessibility. But such initiatives remain marginal to mainstream framework development. The people most affected by AI systems—gig workers subject to algorithmic management, students surveilled by proctoring software, communities targeted by predictive policing—rarely participate in defining the literacy supposedly meant to empower them.

This exclusion matters because literacy frameworks encode assumptions about agency. When [14] warns about AI-generated CSAM, it frames literacy as defensive—teaching children to recognize threats. But what if literacy also meant understanding how platform design enables such harms, or organizing to demand different architectures? The narrow definition of literacy as individual skill forecloses collective responses.

[14] Images et vidéos pédocriminelles crédibles générées par IA se multiplient, alerte l'IWF

Beyond Individual Mastery: Collective Dimensions of Literacy

The individualistic focus of most AI literacy frameworks—teaching each person to navigate AI systems—misses how AI challenges require collective responses. [13] demonstrates this vividly, showing how coordinated AI systems can manufacture false consensus that no individual's critical thinking can counter alone. The solution isn't better individual detection skills but what researchers call "collective sensemaking infrastructure."

[13] How AI swarms manipulate public opinion and how to fend them off

Some frameworks gesture toward collective dimensions. [7] found that students who learned AI concepts in groups developed more sophisticated critical perspectives than those studying individually. The social process of questioning, debating, and collaborative exploration proved essential for moving beyond surface understanding. Yet most

[7] Bridging AI Literacy and Technology Adoption Among Students

frameworks still assess literacy through individual competencies rather than collective capacities.

[15] illustrates why individual approaches fail: AI companions that manipulate children work precisely by isolating them from social support networks. Teaching individual children to recognize manipulation is less effective than creating peer networks where experiences can be shared and validated. The literacy needed here is fundamentally social.

[15] Les chatbots dotés d'intelligence artificielle augmentent le risque d'exposition des enfants à des contenus préjudiciables

The Stanford framework partially recognizes this with its "pedagogical" dimension in [22], acknowledging that teaching others deepens understanding. But it stops short of recognizing literacy as irreducibly collective—not just individuals who happen to share knowledge, but communities that develop shared frameworks for understanding and response.

[22] Understanding AI Literacy | Teaching Commons

Indigenous data sovereignty movements offer an alternative model. Rather than focusing on individual data literacy, these movements assert collective rights over community information, developing governance frameworks that no individual could implement alone. Applied to AI literacy, this might mean neighborhood groups creating shared standards for acceptable AI use, or student bodies collectively negotiating with administrations about surveillance technologies.

The European Schools framework [18] edges toward this recognition with its emphasis on "collective governance," but implementation remains largely individualized. True collective literacy would require new pedagogies—teaching groups to investigate AI systems together, make collective decisions about adoption, and organize responses to harmful deployments.

[18] PDF Cadre pour l'utilisation pédagogique de l'intelligence artificielle...

Toward Participatory AI Literacy

What would AI literacy look like if designed not for compliance or adaptation but for democratic participation? The evidence points toward several principles that current frameworks largely miss. First, literacy must be understood as the capacity to shape AI systems, not merely navigate them. This means understanding not just how AI works but how it's made—the labor conditions, business models, and design decisions that produce particular configurations.

Second, effective AI literacy recognizes the epistemological crisis as central, not peripheral. Rather than offering technical fixes, it develops what [8] calls "epistemic resilience"—communities' collective capacity to establish truth through social processes when technical markers fail.

[8] Deepfakes and the crisis of knowing - UNESCO

This isn't about better deepfake detection but about strengthening the social infrastructures through which knowledge circulates and stabilizes.

Third, participatory literacy acknowledges that different communities need different literacies. [eSafety report shows AI companions are putting children at...] reveals risks specific to young people that adult-focused frameworks miss entirely. Similarly, workers subject to algorithmic management need literacies focused on collective bargaining and regulatory frameworks, not individual optimization. One-size-fits-all approaches reproduce exclusions.

Fourth, the boundaries between AI literacy and other literacies must be understood as porous. [23] demonstrates that "AI misinformation" is inseparable from broader information ecosystems and political economies. Teaching AI literacy in isolation from media literacy, data literacy, and civic education creates artificial boundaries that hinder understanding.

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

Finally, participatory AI literacy is necessarily political. It can't remain neutral about concentrations of power in tech companies, the use of AI for surveillance and control, or the exploitation embedded in training data. Frameworks that avoid these politics in the name of objectivity aren't neutral—they're actively supporting the status quo. As [9] asks, the question isn't whether AI will be governed but how and by whom.

[9] Does A.I. Need a Constitution?

The path forward requires more than adding critical units to existing frameworks or expanding individual competencies. It demands reimagining literacy as collective capacity for democratic participation in shaping our technological futures. This isn't about training better users or even more critical consumers, but about cultivating citizens capable of saying not just how AI should be used, but whether it should be built at all, and if so, by whom and for what purposes.

The battle over AI literacy matters because it's ultimately a battle over agency in an automated world. Current frameworks, despite their sophistication, too often prepare people to adapt to AI systems designed elsewhere for purposes they don't control. What we need instead are literacies of resistance and reconstruction—ways of understanding AI that open possibilities for different configurations, different values, different futures. The question isn't what skills people need for the AI age, but what collective capacities we need to ensure the AI age serves human flourishing. That's the literacy worth fighting for.

References

1. AI & Data Competencies: Scaffolding holistic AI literacy in Higher...
2. AI adoption in schools global trends, benefits, and challenges
3. AI Governance in Higher Education: Policy, Ethics & Risk
4. AI Literacy in K-12 and Higher Education in the Wake of Generative AI...
5. AI Quests from Google teaches AI literacy to kids
6. Anthropic Economic Index report: Learning curves
7. Bridging AI Literacy and Technology Adoption Among Students
8. Deepfakes and the crisis of knowing - UNESCO
9. Does A.I. Need a Constitution?
10. Empowering Learners for the Age of AI
11. Ethics and governance of generative AI in education: a systematic...
12. Gouvernance Algorithmique et Écosystèmes Éducatifs : Vers une Triade...
13. How AI swarms manipulate public opinion and how to fend them off
14. Images et vidéos pédocriminelles crédibles générées par IA se multiplient, alerte l'IWF
15. Les chatbots dotés d'intelligence artificielle augmentent le risque d'exposition des enfants à des contenus préjudiciables
16. Making AI work for schools - Brookings
17. ONU : Les risques croissants des contenus IA pour les enfants
18. PDF Cadre pour l'utilisation pédagogique de l'intelligence artificielle...
19. PDF L utilisation pédagogique, éthique et légale de l intelligence...
20. Prompt Engineering for Education - Substack
21. Towards an AI-Literate Future: A Systematic Literature Review - Springer
22. Understanding AI Literacy | Teaching Commons
23. We Looked at 78 Election Deepfakes. Political Misinformation Is Not an...