

# **The Continued Failure in AI Literacy: AILit produced a starting point halfway through the race and called theory a framework**

Basil C. Puglisi, MPA  
*A Human-AI Collaboration*

In June 2026, the OECD and the European Commission released *Empowering Learners for the Age of AI*, an AI literacy framework for primary and secondary education developed with CodeAI, formerly Code.org, and a group of international experts (OECD & European Commission, 2026). It carries the standing to shape how a generation learns to work with artificial intelligence, and it feeds the PISA 2029 assessment of media and AI literacy. The document is careful and well sourced, organized around four domains of learner competence: Engage, Create, Manage, and Shape AI. It arrives in 2026, late to its own subject, consolidating material that was on the record more than two years earlier and presented in draft along the way. What it delivers, after that long gestation, is a theory wearing the title of a framework.

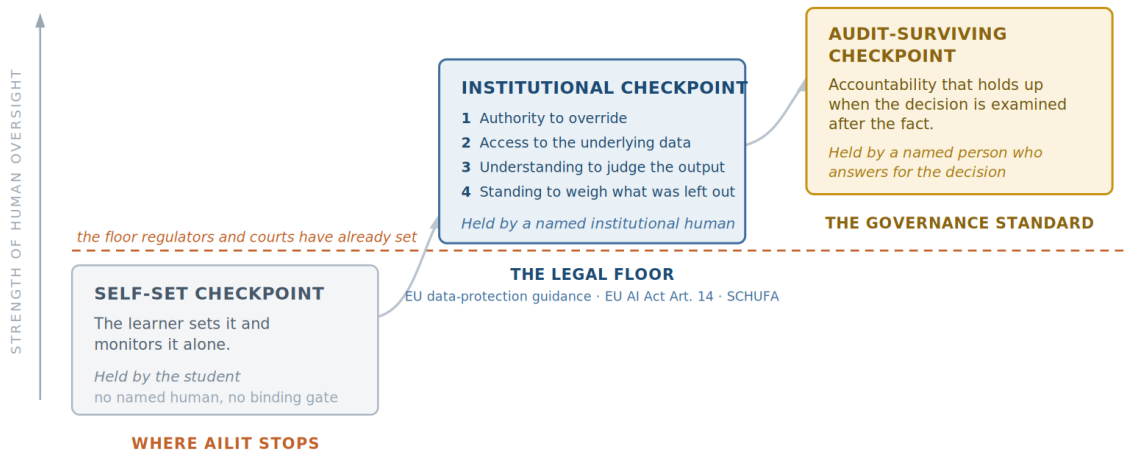
The diagnosis is sound, and the data behind it is real. Eurostat found that 63.8% of 16-to-24-year-olds in the EU used generative AI in 2025, roughly twice the rate of the general population (Eurostat, 2026), while the OECD's own teaching survey reports that around one in three teachers have used AI in their work and three in four say they lack the knowledge or skills to teach with it (OECD, 2025). The need is established, and the urgency is documented. The framework deserves credit for naming both without hedging, and that is where the credit ends.

A diagnosis is not a method, and this is where the document stops. It names the competences a student should hold and supplies almost nothing for how a student builds them. The reason is that the framework sets aside the one body of research that bears directly on the question. Sixty years of learning science locates durable learning in the method of use, in friction deliberately built into the process, in scaffolding that is withdrawn so the learner internalizes the capability, and in the metacognitive monitoring that makes cognitive work observable, none of which appears as a requirement anywhere in the four domains. The framework gestures toward that research for authority and then declines to require what it prescribes. The proven account of how cognitive development happens is on the record, and the framework does not operationalize it.

The path forward is the one the framework declines to take. It is to stop hedging and to operationalize, which means treating the student's cognitive development as the goal, building the friction that produces it, measuring whether that friction is working, and placing a named human at a binding checkpoint who answers for each consequential decision about AI use. Method governance and a named accountable human at every level are what make AI literacy more than exposure to the tools. This analysis takes them in turn: the diagnosis the framework gets right, the method it leaves out, the reason that gap persists, and the requirement it should have set.

## Three definitions of a checkpoint

AILit teaches the self-set checkpoint, while the standard that actually binds was already defined in law.



A learner's self-set milestone sits beneath even the rubber-stamp review the courts have already rejected as insufficient.

Figure 1. The three definitions of a checkpoint, from the learner's self-set version that AILit teaches to the institutional and audit-surviving checkpoints the field already requires.

## A starting point, arriving late

The framework's own promotion claims more than the framework delivers. The OECD's education channel announced that, "instead of preparing students to adapt to AI, we empowered them to shape how it evolves" (OECD Education, 2026). The Shape AI domain is where that promise would have to live. The domain contains classroom design exercises: learners study model cards and propose improvements to AI systems through exemplars. There is no binding authority anywhere in the domain, no mechanism by which a student shapes how a deployed system evolves. The promotion offers agency over the trajectory of the technology; the architecture supplies a worksheet.

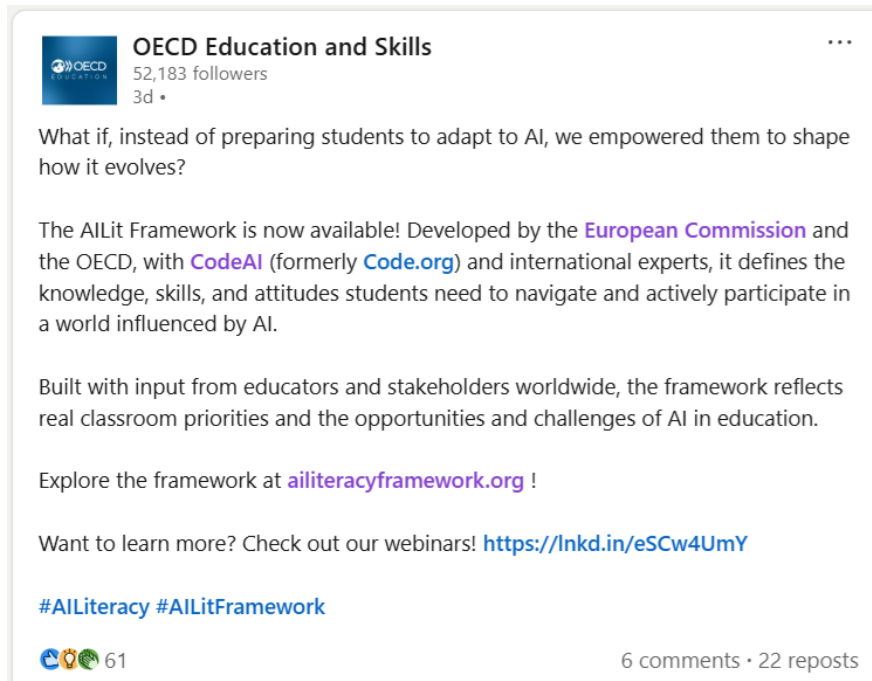


Figure 2. The OECD Education and Skills LinkedIn post promoting the AILit Framework (OECD Education, 2026).

The framework describes itself as "a starting point for establishing a common language," and its own definition of AI literacy "builds on existing definitions" drawn from instruments already two years old (OECD & European Commission, 2026). A consolidation of 2024 materials, published in 2026 as a beginning, arrives at a moment when the operational layer it omits was already specified and dated. The starting line it proposes to mark had been crossed.

## A taxonomy of users in a world that needs systems

Credit the framework for what it is. As a taxonomy of learner skills, it is competent, readable, and usable by teachers who need a shared vocabulary. The charge here is that it aims at the wrong target, not that the work is poor.

Its ceiling shows in two places. The framework builds competent users of a single AI system, and the accountable system that surrounds deployment stays outside its scope. It is also single-model by assumption in a world that has moved past one. Nothing in the framework recognizes that triangulating a question across several AI systems is what manufactures the dissent and source diversity a single model cannot produce on its own. The architecture of friction, the part that makes disagreement visible and forces a human to arbitrate it, falls outside the design before governance is ever reached.

## Friction is structural, not a skill

The reframe that exposes the short-sightedness is this: friction is a property of the system, not a trait of the learner. Decades of cognitive science establish that the *method* of use, rather than the tool, produces the learning. Vygotsky (1978) located learning in the distance between what a learner manages alone and what they manage with help, and Wood, Bruner, and Ross (1976) named the instructional method that operationalizes it: scaffolding, support that is calibrated to a learner's reach and then deliberately withdrawn so the learner internalizes the capability. Chi and Wylie (2014) ordered learning by mode of engagement, with interactive and constructive work producing more durable gains than passive reception.

Learners build durable capability by working through difficulty. Festinger (1957) located reasoning in the productive discomfort of holding a contradiction. Risko and Gilbert (2016) drew the line between strategic offloading, which frees cognitive resources for the hard part of a task, and the suboptimal kind, which substitutes the tool for the thinking. Which one occurs depends on the structure surrounding the task.

The framework develops judgment in the learner. The governance frontier installs accountability in the system, and it manufactures friction across models. ALLit reaches the first and stops before the second.

## The doorway, and the refusal to cross it

The center of the argument is that ALLit reaches the vocabulary of governance and empties it. Its Manage AI domain is defined as the instruction to "divide work intentionally between humans and AI" (OECD & European Commission, 2026). That is allocation, not authority. Even where the framework uses the right words, it houses them inside the learner's own self-monitoring. At the advanced level, learners "establish checkpoints" and "monitor progress against success criteria"; at the intermediate level, they consider "when human judgement should override the algorithm"; at the basic level, they make decisions that "support accountability, learning and fairness" (OECD & European Commission, 2026). In every instance the checkpoint is a study habit. There is no named institutional human, no binding gate, and no accountability that survives an audit.

The standard the framework falls short of is not aspirational. Regulators and courts have already had to define what a genuine human checkpoint is, precisely because token oversight was being passed off as the real thing. Under European data-protection law, the guidance endorsed by the European Data Protection Board requires that oversight of an automated decision be meaningful rather than a token gesture, carried out by someone who has the authority and competence to change the decision and who undertakes a thorough assessment of all the relevant data (Article 29 Working Party, 2018). The European Union's AI Act carries the same logic into its rules for high-risk systems, requiring that the assigned human understand the system well enough to interpret its output and decide against it (Regulation (EU) 2024/1689, Article 14). Read together, the field converges on roughly four conditions for a checkpoint that counts: authority to override, access to the underlying data, enough understanding of the system to judge its output, and the standing to weigh what the system left out. The Court of Justice of the European Union reinforced the floor in its SCHUFA ruling, holding that a decision does not escape accountability because a person formally signs off on it; what matters is the whole situation and whether the automated step plays a determining role (Court of Justice of the European Union, 2023; see also Regulation (EU) 2016/679, Article 22).

That body of law governs automated decisions with legal or significant effects, and it is invoked here for one purpose only: it is the place where the field has already written down what a real checkpoint looks like. The claim is not that a school curriculum violates the GDPR. The claim is narrower and harder to answer. The field already defines a genuine checkpoint, ALLit's version meets none of the four conditions and names no human at all, and the framework neither reaches that standard nor refers to it. A learner's self-set milestone sits beneath even the rubber-stamp review that the courts have already rejected as insufficient. The standard is borrowed as a floor, not a ceiling. If a credit or insurance decision must clear that bar before it can bind a person, a framework preparing a generation for AI-mediated work has no good reason to teach a definition of human oversight that sits below it.

### **The call it would not make**

A framework can fail to build the checkpoint and still show leadership by demanding that someone build it. ALLit does not make that demand either. Its guidance for education policymakers gestures that policy should "uphold governance, ethics and responsible use" (OECD & European Commission, 2026), and it stops there. Its organizing principle is that "AI literacy is a shared responsibility" (OECD & European Commission, 2026). Responsibility spread across everyone is accountability assigned to no one; the phrase diffuses the very thing governance exists to locate. The systematic literature on AI governance is consistent on this point. Across the field, the recurring distinction is between principles that converge easily and the implementation that does not follow (Jobin et al., 2019), and a governance instrument is judged by whether it answers who is accountable, for what, when, and how (Batool et al., 2025). ALLit answers what is to be governed and leaves the accountable party unnamed and undemanded.

The contrast sharpens against the European Union's own binding law. Since February 2025, Article 4 of the EU AI Act has required providers and deployers of AI systems to ensure a sufficient level of AI literacy among their staff and those acting on their behalf, and Article 26 requires that the people assigned to oversee high-risk systems hold the competence, the authority, and the support to do it (Regulation (EU) 2024/1689, Articles 4 and 26). That obligation is binding, it attaches to named parties, and it ties literacy to accountability for deployment. The ALLit framework, co-authored by the European Commission, addresses the same subject for the next generation and makes none of it required. That duty has applied since February 2025, and its enforcement phase opens in August 2026. A proposed amendment would soften the obligation from ensuring a level of literacy to supporting its development. The law already treats AI literacy as a duty owed by those who deploy AI; the framework treats it as a competence hoped for in those who learn.

### **The answer was already on the record**

The operational layer ALLit omits was available to it. The measurement instruments discussed in the next section were published in 2025, the experimental evidence that structured prompting repairs cognitive offloading was published in October 2025, and method-governance approaches built around a named human at a binding checkpoint were public and dated through 2025. The author's own contemporaneous analysis documented this exact failure class, AI treated as a tool with the methods of its use left out as the variable that decides the cognitive

outcome, roughly six weeks before AILit was finalized. The point is precedence, not authorship. The answer existed and was on the record, and the case that AILit fell into the gap rests on the external research cited throughout this analysis rather than on any single prior document. What the record removes is the defense that the operational layer had not yet been articulated. It had.

## **Measurement, deferred while the movement is already underway**

A framework reveals what it believes by what it agrees to measure, and by when it agrees to measure it. The assessment that AILit feeds, the PISA 2029 Media and Artificial Intelligence Literacy framework, is deferred to a single cohort-level instrument whose first results are not expected until 2031 (OECD, 2026b). It contains no measure of whether the prescribed engagement actually happened during any given interaction, and it carries no named accountability. The seam where measurement should bite is the hedge rendered in instrument form.

The deferral would be defensible if the measurement question were genuinely open. It is not. Quantitative measurement of how individuals work with AI is already an active, peer-reviewed area of research, and the instruments predate AILit's publication. Ganuthula and Balaraman (2025) set out a multi-dimensional framework for measuring an individual's collaboration with AI systems, treating the human-AI pairing itself as the unit of measurement. Sidra and Mason (2025) validated scales that measure AI-specific literacy and the metacognition a learner exercises during collaboration, distinguishing it from general metacognition. Both appeared in 2025, before AILit was finalized. The field has begun to measure the thing AILit defers, at the level of the individual and the interaction, well ahead of any population snapshot.

This is the same conversion applied to measurement. The requirement to measure use where it happens becomes a deferral to a distant assessment, and a framework that genuinely intended its prescribed engagement to occur would measure whether it occurred. The instruments that exist are early, and none is yet validated for an assessment the size of PISA. They exist nonetheless, they are entering practitioner use, and they measure at the level of the individual and the interaction that a population snapshot cannot reach. One worked example comes from the author's own practice: the Human Enhancement Quotient, with its Augmented Intelligence Score, a behavioral instrument that decomposes augmented intelligence into measurable dimensions and reports a composite, applied and dated across 2025 and 2026 and offered here as illustration rather than as evidence. The measurement question is no longer blank, and AILit deferred it anyway. The point compounds the charge of arriving late, because the framework is behind the field on measurement as well as on governance.

## **The conviction gap, and the crisis beneath it**

The prose will not stake itself. The evidence is "still mixed"; "further research is needed"; adaptation of the framework is "both expected and encouraged"; teachers are "encouraged to draw from their pedagogical expertise"; the culminating domain "may seem new or unfamiliar" (OECD & European Commission, 2026). A framework that believed its own load-bearing claim, that active and scaffolded engagement produces durable learning where passive tool use does not, would mandate the structure that guarantees the engagement. It describes the destination and makes every step toward it optional.

The asymmetry that gives this away is specific. AILit does not cite the cognitive-science foundations and then hedge its own prescriptions; it does not cite those foundations at all. The names that anchor the science of scaffolding, engagement, and metacognition appear nowhere in its bibliography. The real asymmetry is that where AILit cites an empirical *risk* finding, it states the finding firmly and frames the remedy as optional. It reports, accurately, that offloading effort to AI "may foster 'metacognitive laziness,' weakening learners' critical thinking and self-regulation" (Fan et al., 2025). It then requires none of the structure that would counter the risk it just named. It even concedes the core point on its own, observing that learners using AI "may indeed complete higher quality work" yet "these outcomes do not translate into durable learning gains" (OECD, 2026a), and still asks nothing of the practice that would close that gap. The harm is stated as fact; the fix is offered as a suggestion.

Beneath the hedge sits the crisis that explains it. AILit harvests each source for its authority and strips the part that would create an obligation, and the pattern shows in two forms. The first is the use of a term without its source. The framework speaks of "scaffolding," "metacognition," and "active engagement," and sources these to its own policy documents rather than to the research that defines them. Scaffolding, grounded in Vygotsky's zone of proximal development and named by Wood, Bruner, and Ross, is support that is deliberately withdrawn so the learner takes over; AILit uses the word as a verb for what teachers do to AI use, and the withdrawal, the obligating part, is gone. The same holds for metacognition, which Flavell (1979) defined as the monitoring of one's own thinking; the framework lists the term and omits the discipline the source describes. The second form is the use of half of a source. AILit cites the alarm and omits the fix. It draws on Gerlich's correlational finding that AI use is associated with diminished critical thinking (Gerlich, 2025a), and it omits the same author's experimental finding, available before AILit was finalized, that structured prompting reduces offloading, improves critical reasoning, and functions as a low-cost governance tool (Gerlich, 2025b). One author, one body of work, and the framework keeps the warning while dropping the remedy, including the part of the remedy that names itself governance.

The decontextualization is selective, and it runs in one direction every time. Context is retained where it is safe and removed where it would bind. That selectivity is the proof that the optionality is a choice. A framework carrying scaffolding with its withdrawal, Gerlich with his experiment, and Fan with the implication of his own finding could not in good conscience leave the structure optional. The missing context is the hinge between a framework that is behind the evidence and one that will not commit to it.

Two cautions keep this charge honest. Where AILit does cite an empirical finding, it represents it accurately, and it makes the productivity-versus-learning concession itself; the charge is the selective stripping of obligation, not poor sourcing. And the wider literature on the same question is itself genuinely mixed: the same model can lift practice performance while leaving later, unaided performance worse (Bastani et al., 2025), and a meta-analysis of more than a hundred studies found that human-AI combinations, on average, underperformed the better of human or AI alone (Vaccaro et al., 2024). The evidence does warrant caution. It does not warrant a framework that names the structure and then declines to require it.

## **Why this happens, and why it matters at scale**

Institutions default to diagnosis because diagnosis is safe and operationalization is exposed. Naming a problem costs nothing; specifying who must hold the checkpoint, and who answers when it fails, invites the disagreement and the liability that an institution is built to avoid. The

result is a framework that describes the destination at length and requires nothing on the way to it.

The same stripping of the method layer is visible in the opposite direction, and the author has documented it at length in a companion case study (Puglisi, 2026). In January 2026, a credentialed neuroscientist told the United States Senate Commerce Committee that classroom technology harms learning because human cognition is biological and screens circumvent it, and he ruled out method governance by name, attributing the failure to evolved human biology that no method of deployment could change. The claim traveled as a viral clip viewed more than two million times, and it now appears in state and federal legislation, including a Missouri bill that passed its House and the federal Kids Off Social Media Act the hearing was built to support (Mullin, 2026).

The valence is the inverse of ALLit's, since that testimony blames the tool where ALLit embraces it. The omission is identical. Neither installs a named human accountable for the method, and both treat the medium as the variable that settles the outcome when the research locates the outcome in the method. The reach differs in ALLit's direction, one viral hearing against a framework that feeds an international assessment, so the same omission travels further inside a standard that shapes national curricula than inside one witness's testimony.

Two further amplifiers sit inside the framework itself. The first is that ALLit does not merely omit the checkpoint; it redefines the word downward into self-monitoring and teaches that redefinition as literacy, at continental scale, into PISA 2029. A generation taught that a self-set milestone is a checkpoint is a generation inoculated against recognizing the real thing when its absence matters. The second is the framework's own evidence turned back on it. A document whose entire subject is the sound use of AI commits, in its own construction, the precise failure described above: it cites the friction research and then declines to require friction, keeping Gerlich's alarm and dropping Gerlich's fix. The instruction the framework gives to students, to engage actively and reflect on their use, is the instruction the framework does not follow in its own treatment of the evidence.

The pattern is not confined to one continent or one institution. In the same season, the United States Department of Labor published its own AI literacy framework as voluntary guidance, built on five content areas, sequencing the evaluation of AI outputs fourth and asking workers to "ensure accountability for outcomes" (U.S. Department of Labor, 2026). It too names accountability and assigns it to no one in particular, with no named human, no binding checkpoint, and no mechanism. A national workforce agency and a transnational education body, working from different mandates on different continents, reached the same gap in the same year, which points to a structural pattern rather than a quirk of one document. The scope difference only sharpens it. The Department of Labor framework is voluntary employer guidance, while ALLit feeds an international assessment that shapes national curricula, so the identical omission carries further from the OECD and the European Commission than it does from the Department of Labor.

## **The names are borrowed from a broken gate**

The framework arrived with a sequence of launch webinars whose speakers are members of the expert group that wrote it, drawn from the OECD, the European Commission, CodeAI, and a roster of invited specialists. They convened to present the framework and to walk educators through putting it into practice, and the recording circulates afterward as validation that the

approach is sound. A group convened to present its own work can only perform consensus, and the recording carries that performance outward as proof.

The readability the document earns, and the authority the webinars project, rest on credentialed names and on two institutions whose endorsement is meant to settle the matter. That borrowed credibility is what lets a theory pass for a framework. The names give it the appearance of validated knowledge. The difficulty is that the system those names come from is already compromised in two documented ways, both older than generative AI and neither corrected by this exercise.

The first is the measurement foundation. The cross-population cognitive assessments that education treats as ground truth, the same family of instruments AILit feeds through PISA, were built and validated on a narrow band of humanity. Henrich, Heine, and Norenzayan (2010) established the problem: the Western, Educated, Industrialized, Rich, and Democratic populations on which the behavioral sciences built their findings are systematically unrepresentative, and conclusions drawn from them do not generalize. The critique did not resolve in 2010. Atari and colleagues (2023) documented that roughly ninety percent of psychological research samples still draw from populations that constitute about twelve percent of humanity, and that the findings do not replicate cleanly elsewhere. The author's companion case study traces how this same measurement problem shaped education and its data well before AI entered the classroom (Puglisi, 2026). The instruments a generation will be measured against carry a validity problem the field has known about for fifteen years and has not fixed.

The second is the gate itself. Peer review, the mechanism that is supposed to certify which claims are sound, fails at the most basic check. A Lancet audit found fabricated citations in the biomedical literature rising roughly tenfold between 2023 and early 2026, reaching one in 277 papers, with 98.4 percent of the flagged papers still uncorrected when the audit closed (Topaz et al., 2026). They survive because the gate does not look. In one large survey, 76.7 percent of peer reviewers reported that they do not thoroughly check references, and 74.5 percent judged peer review ineffective at catching citation errors (Xu et al., 2026). The sharpest case sits at the top of the field, where one hundred fabricated citations spread across fifty-three accepted papers at a leading machine-learning conference passed review by three to five expert reviewers each (Ansari, 2026). Approved, peer-reviewed work now carries hallucinated sources and invented data, which is direct evidence that the gatekeeping a credential rests on is not performing the verification its authority claims.

Put the two together and the conclusion follows: AI made an existing failure visible. The gate was already open before the first hallucinated citation, and the measurement foundation beneath it was already unrepresentative. The education system that deployed unvalidated technology to children, built its assessments on a narrow population, left peer review unable to catch a fabricated citation, and named no one to answer for any of it is the same system now convening its own experts to certify its own AI literacy theory. A gate that cannot guard its own record cannot certify the repair of what it failed to guard. This is the structural reason the missing piece matters. Closing the gap takes an external checkpoint and a named human who answers for the result, the one thing the framework declines to require and the one thing its own validation process does not supply.

## The call ALLit should have made

What leadership would have looked like is not a mystery, because the standard already exists. The framework calls itself a starting point, and a starting point is a fair thing to be, yet a starting point that leaves every step optional becomes a permission slip for the status quo. Requiring structure would not have exceeded the framework's remit. It could have required, without overstepping, that every AI-assisted assignment carry a named human checkpoint and a recorded account of that check. That is a structural requirement an international framework can set, and it is a different matter from prescribing any particular classroom method. A framework willing to stake itself would have named the human who holds the checkpoint and required the gate, using the meaningful-human-involvement standard the field has already written into law as the floor. That gate is not met by a human-in-the-loop placed only to review what the system produced, because review without ownership is the rubber stamp the courts have already rejected. The requirement's center is a named person who owns the decision and answers for it, with the authority to override, the access to verify, the understanding to judge, and the accountability that survives an audit. Methods that operationalize that standard already exist, the author's own Checkpoint-Based Governance among them, which is what makes ALLit's silence a decision: the requirement was buildable, and the framework chose to leave it optional.

ALLit had the evidence, the standing, and a finalization date late enough to use both. It produced a careful taxonomy that requires nothing. The framework is the de facto international standard now, and it tells a generation that using AI well is a matter of personal competence rather than governed deployment. The cost of that choice will not be visible in a cohort-level score in 2031. It will be visible in the first generation of workers who were taught that a checkpoint is something you set for yourself, arriving in workplaces where the question is who answers when the AI output ships and is wrong.

The evidence was on the record before the framework was finalized, and the standard for a real checkpoint was already written into law. A framework with this reach could have required it. Choosing to require nothing, with both the evidence and the standard in hand, was itself a statement of what the framework believed about its own content.

## Sources

Ansari, S. (2026). *Compound deception in elite peer review: A failure mode taxonomy of 100 fabricated citations at NeurIPS 2025* (arXiv:2602.05930). arXiv. <https://arxiv.org/abs/2602.05930>

Article 29 Working Party. (2018). *Guidelines on automated individual decision-making and profiling for the purposes of Regulation 2016/679* (WP251rev.01). European Commission. (Endorsed by the European Data Protection Board, 25 May 2018.)

Atari, M., Xue, M. J., Park, P. S., Blasi, D. E., & Henrich, J. (2023). *Which humans?* (PsyArXiv preprint). <https://doi.org/10.31234/osf.io/5b26t>

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, *122*(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>

Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: A systematic literature review. *AI and Ethics*, *5*, 3265–3279. <https://doi.org/10.1007/s43681-024-00653-w>

Chi, M. T. H., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, *49*(4), 219–243. <https://doi.org/10.1080/00461520.2014.965823>

Court of Justice of the European Union. (2023). *SCHUFA Holding (Scoring)*, Case C-634/21, judgment of 7 December 2023, ECLI:EU:C:2023:957.

Eurostat. (2026, February 10). *64% of 16–24-year-olds used AI in 2025*. <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/edn-20260210-1>

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. *British Journal of Educational Technology*, *56*, 489–530. <https://doi.org/10.1111/bjet.13544>

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford University Press.

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*(10), 906–911. <https://doi.org/10.1037/0003-066X.34.10.906>

Ganuthula, V. R. R., & Balaraman, K. K. (2025). Artificial intelligence quotient framework for measuring human collaboration with artificial intelligence. *Discover Artificial Intelligence*, *5*, Article 268. <https://doi.org/10.1007/s44163-025-00516-1>

Gerlich, M. (2025a). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, *15*(1), Article 6. <https://doi.org/10.3390/soc15010006>

Gerlich, M. (2025b). From offloading to engagement: An experimental study on structured prompting and critical reasoning with generative AI. *Data*, *10*(11), Article 172. <https://doi.org/10.3390/data10110172>

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*(2–3), 61–83. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1601785](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1601785)

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>

Mullin, J. (2026, January 16). *Congress wants to hand your parenting to Big Tech*. Electronic Frontier Foundation. <https://www.eff.org/deeplinks/2026/01/congress-wants-hand-your-parenting-big-tech>

OECD. (2025). *Results from TALIS 2024: The state of teaching*. OECD Publishing.

OECD. (2026a). *OECD Digital Education Outlook 2026*. OECD Publishing. <https://doi.org/10.1787/062a7394-en>

OECD. (2026b). *Navigating an evolving digital world: First draft of the Media and Artificial Intelligence Literacy (MAIL) assessment framework (PISA 2029)*. OECD Publishing.

OECD & European Commission. (2026). *Empowering learners for the age of AI: An AI literacy framework for primary and secondary education*. OECD Publishing. <https://doi.org/10.1787/65cd27d4-en>

OECD Education. (2026, June). *The AILit Framework is now available* [Post]. LinkedIn. [https://www.linkedin.com/posts/oecd-education\\_ailiteracy-ailitframework-activity-7473313581699104768-sWPu](https://www.linkedin.com/posts/oecd-education_ailiteracy-ailitframework-activity-7473313581699104768-sWPu)

Puglisi, B. C. (2026). *How credentialed professionals shape policy when method governance is stripped* [The Horvath case study]. BasilPuglisi.com. <https://basilpuglisi.com/how-credentialed-testimony-outpaces-research-horvath-case-study/>

Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation), Article 22.

Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), Articles 4, 14, and 26.

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>

Sidra, S., & Mason, C. (2025). Generative AI in human-AI collaboration: Validation of the Collaborative AI Literacy and Collaborative AI Metacognition Scales for effective use. *International Journal of Human–Computer Interaction*. Advance online publication. <https://doi.org/10.1080/10447318.2025.2543997>

Topaz, M., Roguin, N., Gupta, P., Zhang, Z., & Peltonen, L.-M. (2026). Fabricated citations: An audit across 2.5 million biomedical papers. *The Lancet*, 407(10541), 1779–1781. [https://doi.org/10.1016/S0140-6736\(26\)00603-3](https://doi.org/10.1016/S0140-6736(26)00603-3)

U.S. Department of Labor. (2026). *Training and Employment Notice No. 07-25: The U.S. Department of Labor's Artificial Intelligence Literacy Framework*. <https://www.dol.gov/agencies/eta/advisories/ten-07-25>

Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour*, 8(12), 2293–2303. <https://doi.org/10.1038/s41562-024-02024-1>

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.

Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2), 89–100. <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>

Xu, Z., Qiu, Y., Sun, L., Miao, F., Wu, F., Wang, X., Li, X., Lu, H., Zhang, Z., Hu, Y., Li, J., Luo, J., Zhang, F., Luo, R., Liu, X., Li, Y., & Liu, J. (2026). *GhostCite: A large-scale analysis of citation validity in the age of large language models* (arXiv:2602.06718). arXiv. <https://arxiv.org/abs/2602.06718>

---

**Disclaimer.** The author is not a lawyer, and nothing in this article is legal advice. The GDPR provisions, the SCHUFA ruling, the Article 29 Working Party guidance, and the EU AI Act duties discussed here, including the enforcement timeline and the proposed amendment to Article 4, should be confirmed against the current text and authority, since regulatory instruments and their interpretation can change. Any decision about deploying or governing AI systems should be made with qualified counsel. The author is an independent practitioner and author who may profit in other ways from research and content like this.

**The Other AI: Audio Briefings on Augmented Intelligence and AI Governance**

[Spotify](#); [Apple Podcasts](#); [Amazon Music](#); [YouTube Playlist](#);

#AIassisted using HAIA Ecosystem

---