

AI Literacy in Education across Regulatory Frameworks: A Comparative Analysis of the European Union, United States and China

Ana Zidărescu

Independent researcher, Brussels

ana_zidarescu@yahoo.com, [ORCID](#)

DOI: 10.63467/all17.art14

**The views and opinions expressed in this paper are solely those of the author and are made in a strictly private capacity.*

They do not reflect the official policy or position of the European Parliament, any of its bodies, services, Members, political groups or committees. The author does not act on behalf of, or in the name of, the European Parliament.

Abstract

AI literacy is no longer just an educational objective or an individual competence. State authority now uses it as a governance framework to manage risks and assign accountability for AI use in education. Yet, current literature continues to frame the concept almost entirely around user competence and ethical boundaries. This leaves its growing institutional role, specifically how it acts as a mechanism for legal responsibility, largely unexamined. To address this gap, the article examines how AI literacy is defined, framed, and translated into educational responsibility across three important, yet different governance models: the European Union, the United States and China. These frameworks were selected because they represent three different approaches to AI literacy governance in education: a legal-compliance model in the European Union, a decentralised, and incentive-based model in the United States, and a centralised, state led model in China.

Through a qualitative comparative analysis of legal texts, official guidelines, and education-specific policy documents, the study analyses how the same concept is assigned different regulatory and educational functions across three jurisdictions. The findings prove that AI literacy extends far beyond personal skills and competences, and it now operates primarily as a mechanism

of institutional governance, fundamentally reshaping educational policy, teacher liability, and algorithmic responsibility.

Keywords: AI literacy, AI governance, education policy, responsible AI, AI regulation, European Union, United States, China

1. Introduction

As a concept, AI literacy barely had a moment to settle. When the earliest definitions emerged, between 2016 and 2020, the focus was almost exclusively pedagogical. The entire debate revolved around individual skillset: teaching students to understand basic algorithms, or showing teachers how to actually use these new platforms in their daily teaching (Long & Magerko, 2020; Ng et al., 2021). Yet, the rapid evolution of AI systems has very rapidly outpaced this theoretical focus. Before schools and universities could even properly institutionalise these basic competencies, automated tools flooded academic life. Today, algorithms filter admissions, grade student performance, draft content, or generate lesson planning, often without formal institutional approval (Williamson & Hogan, 2021).

The OECD's *Digital Education Outlook 2026* explicitly flags this exact blind spot. Unlike previous waves of educational technology, the new generation of AI tools is widely accessible for free and frequently used outside any formal procurement process or institutional oversight. Because of this, the current challenge is less about pedagogical preparedness and more about a severe vacuum of authorisation. Whilst early guidelines, such as the *UNESCO Guidance for Generative AI in Education and Research* (2023), established vital baselines for “human-centred use” and “age-appropriate safeguards”, they also highlighted an alarming lack of institutional readiness. Currently, most AI tools shaping the academic work are neither formally approved nor systematically supervised. We are operating in a regulatory grey zone of well intended, but abstract warnings, completely avoiding the most critical administrative question: who is actually accountable when the rules are superficially defined or simply don't exist?

This vacuum forces a fundamental redefinition of AI literacy itself. The concept must cease being viewed merely as a personal pedagogical checklist and instead take on an explicit institutional mandate. Measuring the competency level of students and educators is no longer sufficient. The critical question is now what an academic system actively enforces.

While upskilling efforts remain necessary, educational communities must move beyond basic training and build a tangible “architecture of responsibility” (Holmes et al., 2021). Institutions must determine what acceptable use looks like and decide exactly who has the responsibility when a system fails.

Yet, on a global scale, this push for AI literacy governance intersects with a fragmented geopolitical reality. Recent global events, such as the 2026 India AI Impact Summit, where the US delegation explicitly rejected a unified global AI framework, demonstrate that the major powers remain fiercely loyal to their own models to secure market dominance. By extension, AI literacy also inherits the ideological marks of its parent system and this reality makes a comparative approach a methodological necessity. For this reason, this paper does not treat the European Union, the United States, and China as isolated case studies, but as actual competing administrative philosophies with operational differences. The EU anchors AI literacy directly to legal obligation and compliance standards (European Parliament, 2024). The US deliberately avoids federal mandates, opting instead for a fragmented, decentralized approach (US Department of Education, 2023). Meanwhile, China frames the concept within its centralized administrative machinery (Cyberspace Administration of China, CAC, 2023). The analytical goal of this paper is to understand how each framework justifies its authority at the exact moment AI literacy evolves from personal skillset into a mechanism of institutional governance.

2. Literature Review and Research

The traditional skill-based focus on AI literacy has created a profound disconnect between academic theory and institutional practice. Existing literature remains largely anchored in a pedagogical bubble and existing frameworks, such as the UNESCO *AI Competency Frameworks* (2024), still define AI literacy primarily as a matter of individual knowledge and ethical awareness. This competence-based approach is necessary, but it completely ignores the administrative reality of schools.

Currently, the closest attempt to bridge this gap is the work of Yang et al. (2024), who examine AI policy through the lens of policy diffusion, comparing frameworks across the EU, US, India and China. They explicitly treat national strategies as engines for "domestic talent cultivation" and global competitiveness, they also effectively map the geopolitical race for an AI-ready

workforce, but leave a critical vacuum by not examining AI literacy as a tool of educational governance.

Looking at the broader literature, we observe a similar imbalance. We have an abundance of research on what students and teachers should know, exemplified by the UNESCO AI Competency Frameworks (2024) or Swiecki et al., (2022), but there is little to no documentation on who actually has the administrative responsibility once a system is deployed in a school. As Pesek (2026) notes in *Rethinking Education Governance in the Age of AI*, integrating artificial intelligence in schools shifts who makes the actual decisions.

We also have to acknowledge a massive administrative void when it comes to schools and universities. Unlike the corporate sector, where technology mainly exists to generate profits, in education AI has multiple and sensitive roles. It is both a tool for learning and a subject of study, yet it also evaluates students, filters admissions, and helps in daily administrative tasks. This specific dynamic turns education into a major vulnerability for AI governance (Selwyn et al., 2022). When a single technology can dictate institutional access and grade human performance, it can, ultimately, influence a student's entire academic trajectory. As Panagopoulou, Parpoula and Karpouzis (2025) note, this raises severe issues of equity and institutional accountability. In this context, AI literacy can no longer be just a personal pedagogical checklist. Instead, this reality transforms students and educators from simple users, to subjects of an automated bureaucracy.

To directly address this administrative void, this paper is organised around three research objectives. First, it seeks to redefine AI literacy by reframing it from a strictly institutional perspective, moving the definition from “pedagogical skill” to “governance tool”. Then, it maps and compares how three major political actors, the EU, the US and China, allocate administrative responsibility when regulating AI in their educational systems. On that basis, by analysing these three regulatory models, the paper introduces a practical analytical instrument, the Policy-to-Practice Translation Matrix, to measure the degree to which AI literacy is translated from high-level policy into on-the-ground governance in education.

3. Methodology

Mapping the transition of AI literacy from a policy ambition into an enforceable administrative reality demands more than a standard thematic review. Merely scanning these national (the U.S

and China) and regional (the EU) strategies and searching for keywords like “ethics” or “safety” fails to explain how accountability actually works on the ground and how it is implemented by educational actors in schools and universities (Rizvi & Lingard, 2010; Curaj et al., 2020). Policymakers routinely rely on this type of vocabulary to project a sense of regulatory control, yet this rhetoric is not enough to become an actual administrative tool, and it almost never translates into clear instructions for school boards. Simply counting the frequency of “safe AI” or “responsible AI” in a government document cannot expose an actual chain of command. Ultimately, a keyword-driven approach deliberately avoids the single most critical bureaucratic question a policy must answer: who actually carries the legal responsibility when a system fails while being used in an academic setting?

To answer this exact question, the paper proposes a comparative Critical Policy Analysis (CPA) framework (Diem et al., 2014; Taylor, 1997). Rather than evaluating pedagogical outcomes, the methodology introduces a strict institutional lens that examines official texts through precise governance indicators. To structure this, the paper develops the Policy-to-Practice Translation Matrix, tracking key elements like policy drivers, legal responsibility, and implementation mechanisms across jurisdictions. To apply this matrix, the paper relies on a strict defined corpus of policy documents, spanning from 2022 to April 2026, chosen precisely because they reflect contrasting governance logics.

To map the European Union’s approach, the paper examines texts that strongly anchor educational technology (EdTech) within a rigid risk classification. The primary text examined is the *Artificial Intelligence Act* (European Parliament, 2024), specifically Article 4 and Recital 56, which directly target educational use cases and outline institutional obligations. However, because this legal architecture is actively updating, the corpus also incorporates the proposed *Digital Omnibus on AI* alongside the European Commission’s *Ethical Guidelines*(2026). This addition is critical for tracing how the initial statutory burden is currently being renegotiated, demonstrating how it translates abstract legislative categories into administrative reality.

The United States, meanwhile, actively resists any influence toward a centralised federal mandate. To map this deliberately fragmented approach, the corpus traces the American policy trajectory from its early theoretical principles to its current funding mechanisms. The analysis reviews the White House’s *Blueprint for an AI Bill of Rights* (Office of Science and Technology Policy, OSTP,

2022), the recent *Executive Order 14277* (2025) and the voluntary US Department of Labor *AI Literacy Framework* (2026) against the US Department of Education's funding guidelines. Furthermore, to properly understand the complex administrative reality of the US, the methodology contrasts these soft-law federal documents with state level hard laws, specifically looking at recent legislation like the *Colorado Act* (2024). Together, these documents illustrate a clear strategy to avoid the federal legislative path, highlighting a fragmented governance characteristic that creates major differences between school districts.

Breaking completely from the EU and the US models, China deliberately rejects both the European legal framework and the American free-market and innovation-oriented model, choosing instead to merge compliance with strict ideological control. To understand this third model, the paper extracts governance indicators from the *Interim measures for the management of Generative Artificial Intelligence Services* (Cyberspace Administration of China, CAC, 2023). To track how this strict national security approach actually influences the use of AI in classrooms, the analysis also incorporates the Ministry of Education (MOE) guidance on AI literacy (2024) alongside the new *AI + Education Action Plan* (April 2026).

4. Comparative Analysis and Results

Applying the Critical Policy Analysis (CPA) matrix to the selected texts reveals a major administrative discrepancy. Globally, the public rhetoric on AI literacy is almost identical, promising ethical deployment and student well-being, the actual bureaucratic reality is highly fractured. This geopolitical fracture surfaces the moment we examine the first set of governance indicators: the actual legal weight of the mandate and the precise allocation of institutional responsibility.

4.1 The European Union's approach to AI governance: strictly binding legal architecture

By passing the Artificial Intelligence Act in 2024 as a regulation rather than a directive, Brussels guaranteed the text immediate statutory force across all 27 Member States, entirely bypassing the need for national parliaments to transpose the rules into national law. The defining feature of this legislation is its risk-based approach. Rather than regulating the technology itself, the EU regulates the specific use case, forcing AI systems into a rigid four-tier risk matrix, classified from

“minimal” to “unacceptable“. Under this strict hierarchy that Article 4 of the AI Act changes the rules for academic institutions. Because AI systems used in education end up directly dictating human trajectories, by filtering university admissions, or grading student performances, they are explicitly classified as “high-risk” systems.

Here, AI literacy loses its pedagogical focus and becomes a strict organisational duty. The original statutory text explicitly states that both software providers and institutional deployers, meaning the schools and the universities themselves, must ensure a “sufficient level of AI literacy” to manage these high-risk systems.

Once a university deploys a high-risk educational algorithm, it has to legally guarantee continuous human oversight and maintain incident logs. Also, it cannot shift the liability to the software vendor if an automated grading tool starts hallucinating facts or discriminating. As the statutory deployer, the academic institution carries the immediate administrative responsibility. Selwyn et al. (2022) note that this dynamic turns schools into “automated bureaucracy” hubs, forcing local boards to absorb legal risks for technologies they do not entirely understand. Expecting a local high school to meet the same regulatory standards as a multinational tech company is, however, practically unsustainable. Recognizing this, the European compliance model is currently undergoing a major shift. As negotiations extend into mid-2026, the proposed Digital Omnibus on AI signals a profound administrative shift.

Recognising the severe legal exposure placed on local educators, the Omnibus attempts to simplify and amend this dynamic. The proposed text shifts the risk oversight away from individual deployers, transferring regulatory burden directly to the European Commission and the Member States, under the monitoring of the EU AI Office. This renegotiation shows the EU acknowledging a structural flow in its initial allocation of accountability. By absorbing the compliance burden at state level, the EU is actively redesigning its chain of responsibility in education. Yet, even as this power shift is negotiated, the core European philosophy remains intact: given its strategic role, AI literacy will always require statutory compliance, never a voluntary recommendation.

4.2 The United States: the decentralised, incentive-based model

Unlike Brussels, Washington actively supports a regulatory vacuum at the central level. The federal American model proposes a “soft law” approach, a deliberate strategy of governance

philosophy that privileges market agility and innovation over centralised compliance. As mapped through the White House Executive Order 14277 (2025) and the US Department of Labor AI Literacy Framework (2026), the federal government explicitly frames AI literacy not as a legal duty, but as a workforce capacity-building exercise.

While this decentralised approach gives the corporate sector exactly the agility it wants, for local school boards, this hands-off approach creates a more challenging reality. Because the federal government, through the Department of Education, offers funding incentives and voluntary guidelines rather than a hard legal framework for educators to actually rely on, academic institutions are left to navigate algorithmic responsibility on their own. Also, in the absence of federal rules, individual states have started to act, creating what is essentially a fragmented legislative patchwork. States like Colorado, with its Colorado AI Act passed in 2024, and California have introduced their own hard laws targeting algorithmic discrimination. For a school district, this creates legal inequality. Depending on where they are located, schools are either facing strict state-level penalties or relying entirely on legacy privacy laws, like the 1974 FERPA law, trying to fit old rules to cover modern machine learning. Ultimately, this turns student protection into a “zip code lottery”. A school’s capacity to actually enforce AI literacy ends up depending on its local legislation and tax revenue rather than a guaranteed national standard.

4.3 China: the centralised, state-led model

Breaking completely from both the European compliance framework and the American market-led approach, Beijing treats AI literacy as a strict matter of national security. Looking at the Interim Measures (Cyberspace Administration of China, 2023) and the recent AI + Education Action Plan (April 2026), the state clearly has no intention of letting local schools to assess algorithmic risks by themselves. Instead, the central government handles the entire regulatory audit long before any software reaches the classroom. Before even touching the education market, tech developers must pass severe and rigorous state security assessments to prove that their generated outputs align with “socialist core values”.

For Chinese schools, this creates a strictly top-down reality. Controlled entirely by the Cyberspace Administration of China (CAC), they do not act as legally liable “deployers” as they are in Europe, nor are they left legally exposed like the American schools and universities. They have no

compliance burden because the state takes all responsibility. However, this heavy central filtering alters the very nature of educational technology. Because Beijing controls in such a heavy way all compliance decisions, ideologically unwanted AI never even makes it to classrooms. This completely shifts the concept of AI literacy: it stops being a vehicle for open pedagogical inquiry and turns into a highly monitored instrument for political compliance.

4.4 Cross-regional comparison: the architecture of responsibility

Looking across these three jurisdictions, it becomes clear that their antagonistic positions on AI literacy have almost nothing to do with actual teaching. On paper, their position is practically identical: every region promises to nurture critical thinking, protect student privacy, and prepare learners for an automated economy. But as Ben Williamson (2021) rightly points out, the reality of an educational technology policy is never found in its pedagogical promises, and it's actually located in its administrative wiring. The real geopolitical divide happens the moment we talk about liability: when a system fails, hallucinates, discriminates, or extracts data without consent, who is to be legally held accountable?

The European Union, the United States, and China resolve this crisis in completely different ways. Brussels initially treated AI literacy as a strict institutional compliance, forcing schools and universities to absorb the legal burden of high-risk algorithms, an approach it is now trying to correct through the 2026 Digital Omnibus on AI negotiations. Washington, meanwhile, actively maintains a federal regulatory vacuum to protect market innovation. By refusing to enforce a national framework, the US effectively pushes the administrative burden down to local schools, leaving them legally exposed and financially outmatched by tech developers. Beijing handles the entire risk management right at the top. By treating algorithms as a national security issue, the Chinese state removes any compliance authority from schools, using AI literacy as a vehicle to enforce political conformity.

What these three models show is that we can no longer treat AI literacy globally as just a universal set of technical skills. Depending on the geopolitical context, the concept serves different institutional functions: it operates as a strict legal obligation in the EU, a locally founded workforce competency in the US, and a direct instrument of state control in China.

4.5 The Policy-to-Practice Translation Matrix: a comparative framework

Simply scanning official documents for keywords like “responsible AI” fails to capture how an AI literacy policy actually works in real life. Recent empirical evidence supports this argument. As Park et al. (2026) note, building AI skills without enforcing strict rules makes societies more vulnerable. According to them, AI literacy works safely when supported by mature governance. But while quantitative models confirm that governance is crucial, they cannot explain how these rules actually function once deployed within an academic ecosystem. To bridge this gap, this paper introduces the Policy-to-Practice Translation Matrix, a qualitative analytical tool that maps exactly how abstract regulations translate into daily academic life.

The five indicators comprising this matrix are highly pragmatic. Moving past the surface of official documents, this framework shifts the focus away from what educational policies claim to achieve, asking a much more practical question: who actually has the power and the responsibility to enforce them?

To trace this accountability across these three regions, the matrix applies five specific criteria:

- 1. The policy driver:** Analyzing a policy requires identifying its geopolitical context. This indicator assesses the actual justification for state intervention, whether the goal is to protect human rights, promote market innovation, or secure ideological control.
- 2. Legal responsibility:** This indicator maps who bears the legal consequences when an algorithm fails or discriminates. It also identifies the specific agencies in charge, from the EU AI Office to China’s Cyberspace Administration (CAC).
- 3. Implementation mechanism:** This indicator measures the regulatory weight of the policy, distinguishing between hard law that demands strict compliance and soft-law recommendations.
- 4. Implementation support:** It evaluates whether schools receive the actual financial and administrative resources needed to enforce the rules, or if they are left to absorb the regulatory burden alone.
- 5. Student protection:** This final indicator examines how learners are practically protected in practice. Moving beyond abstract ethical promises, it interrogates the concrete mechanisms in place to prevent an algorithm from causing harm.

Governance indicator	The European Union	The United States	China
The policy driver (<i>Why the state acts</i>)	Risk-based approach. Focuses on human rights and risk management. The main goal is to protect citizens from automated harm.	Market-based approach. Focuses on the economy. The goal is to build a highly skilled workforce and push market innovation.	Security-based approach. Focuses on national security and regime stability. The primary goal is strict ideological control.
Legal responsibility (<i>Who takes the legal blame</i>)	Shifting to macro-level (State / EU). Monitored by the EU AI Office. The burden of ensuring AI literacy is shifting away from schools (as deployers) and companies, moving directly to the European Commission and Member States under the proposed Digital Omnibus	Pushed to micro level (local / district). There is no central AI agency. Schools and universities carry the primary risk under outdated privacy laws, relying entirely on voluntary recommendations from NIST and the Department of Education.	Absolute centralisation. Controlled entirely by the Cyberspace Administration of China (CAC). The state takes all responsibility, and local schools carry zero compliance burden.
Implementation mechanism (<i>How rules are enforced</i>)	Uniform hard law. The AI Act demands strict legal audits based on a four level risk system.	Soft law, decentralised rules. The federal government offers voluntary recommendations. Meanwhile, individual states like Colorado and California are passing their own strict AI laws, creating legal inequality between school districts	Strict state approval. Mandatory state registry. Every AI system must pass a central security assessment before schools can even test them.
Implementation support (<i>What help schools get</i>)	Low / delegated to the deployer. Under the 2024 EU AI Act, schools (as deployers) absorb the compliance burden. While Member States must create oversight agencies, they have no legal obligation to provide direct financial or administrative support for schools, leaving them	Low & highly unequal / delegated to local districts. The federal government provides soft-law guidance and competitive funding, but no guaranteed administrative protection. Support depends entirely on state-level legislation and local tax revenue,	High / centralised state absorption. Schools face zero compliance burden because the central state machinery (CAC) handles the entire regulatory security audit before the software reaches the classrooms. The state absorbs the

	exposed (a regulatory flaw the 2026 Digital Omnibus is now attempting to address)	placing a major influence on the “zip code lottery” where poorer districts are left to manage risks independently.	administrative responsibility completely.
Student protection (Actual safety rules)	Strong legal protection. High / Legally enforced. Algorithms that make decisions about a student’s future are strictly classified as “high-risk”	Vulnerable / outdated legal protection. Left to outdated privacy laws. Schools try to fit the 1974 FERPA law to cover modern machine learning, which is not enough.	State-filtered protection. Subordinated to socialist core values. Protecting the student simply means ensuring the AI aligns with state ideology.

Table 1. The Policy-to-Practice Translation Matrix

As **Table 1** demonstrates, applying these governance indicators reveals that AI literacy is far from a universal educational goal. Instead, it reflects distinct geopolitical agendas. The EU regulates to protect rights, the US prioritizes market innovation, and China restricts to maintain state control. As the matrix shows, the real barrier to safe AI in schools is not the technology itself, but the problem is the mismatch between what governments demand and what they can actually support. For the EU and the US, this gap creates a serious paradox. Both want to dominate the global AI race and expect local schools to train a highly skilled workforce for it. Yet, as the legal liability indicator shows, by pushing the legal risk down to local districts, they trigger a defensive administrative response. Because schools receive fragmented or low implementation support, they are forced to absorb the legal responsibility alone. Historically more cautious and avoiding legal risks (Selwyn, 2021), academic institutions facing the threat of bias lawsuits or data breaches under outdated privacy laws will likely opt for safe and heavily restricted AI tools. Ultimately, institutional self-preservation will sabotage the very pedagogical innovation these policies claim to promote.

In contrast, China resolves this liability right at the top. Because the central government (via the CAC) handles the entire approval mechanism, schools face no legal risk. As the student protection indicator reveals, this administrative safety comes at the cost of pedagogical autonomy. The

technology ceases to be an open platform for learning, and AI literacy is reduced to an instrument for behaviour monitoring and political alignment.

5. Conclusions

This paper began by tracing a fundamental shift in educational technology: AI literacy is no longer just a pedagogical checklist. As academic institutions are flooded with different automated grading tools, content generation, even automated filters for school admission, the challenge has moved from the classrooms directly into institutional bureaucracy. Our objective was to look beyond the pedagogical definitions of AI literacy and to examine how this new “architecture of responsibility” actually translates in reality.

By applying the proposed matrix across the three jurisdictions, the analysis reveals that AI literacy does not exist as a universal governance standard. Instead, it functions as an extension of its parent system and acts the same as any other political instrument does.

More importantly, by placing the three models in competing positions, we discover a major structural flaw: both Washington and Brussels expect academic institutions to cultivate an AI-ready new workforce generation to dominate the global market. Yet, by placing the entire accountability of high-risk systems onto local schools and districts they actually trigger a predictable administrative reaction: institutional self-preservation. On the other hand, China resolves this liability by centralising all approvals at state-level. The cost, here, is that technology serves as a control instrument and a total lack of pedagogical autonomy.

In reality, all three frameworks fail the educational system, because safe AI in education is a matter of institutional responsibility, not just ethics. If governments genuinely want AI-literate students, teachers and overall academic institutions, they must move beyond publishing guidelines or recommendations, and they must also stop forcing schools to act as legal deployers. The education sector needs real legal and financial support to implement a real and effective AI literacy governance. Ultimately, none of these approaches builds AI-ready generations or academic institutions, but they manage to build defensive bureaucracies. Until policymakers learn to protect schools from liability without transforming AI into a tool for authoritarian control, AI literacy will remain a fractured policy ambition.

References

- Colorado General Assembly. (2024). Colorado Artificial Intelligence Act (SB24-205) <https://leg.colorado.gov/bills/sb24-205>.
- Council of Europe. (2026). Draft Recommendation on AI Literacy. Steering Committee for Education (CDEDU) <https://www.coe.int/en/web/education/-/the-steering-committee-for-education-cdedu-approves-a-landmark-framework-on-ai-literacy>.
- Cyberspace Administration of China (CAC). (2022). Administrative Provisions on Algorithm Recommendation in Internet Information Services https://www.cac.gov.cn/2022-01/04/c_1642894606364259.htm.
- Cyberspace Administration of China (CAC). (2023). Interim Measures for the Management of Generative Artificial Intelligence Services https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm.
- Curaj, A., Deca, L., & Pricopie, R. (Eds.). (2020). European Higher Education Area: Challenges for a New Decade. Springer Nature. <https://doi.org/10.1007/978-3-030-56316-5>.
- Diem, S., Young, M. D., Welton, A. D., Mansfield, K. C., & Lee, P. L. (2014). Intellectual landscape of critical policy analysis. *International Journal of Qualitative Studies in Education*, 27(9), pp. 1068-1090.
- European Commission. (2025). AILit Framework - Empowering Learners for the Age of AI: an AI Literacy Framework for Primary and Secondary Education (Review Draft) https://ailiteracyframework.org/wp-content/uploads/2025/05/AILitFramework_ReviewDraft.pdf.
- European Commission. (2026). Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators <https://education.ec.europa.eu/focus-topics/digital-education/actions/plan/ethical-guidelines-for-educators-on-using-artificial-intelligence>.
- European Parliament. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=OJ:L_202401689.

- Holmes, W., Persson, J., Chounta, I. A., Wasson, B., & Dimitrova, V. (2022). Artificial Intelligence and education: A critical review through the lens of human rights, democracy and the rule of law. Council of Europe.
- Long, D., & Magerko, B. (2020) What is AI Literacy? Competencies and Design Considerations. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1-16 <https://dl.acm.org/doi/epdf/10.1145/3313831.3376727>.
- Ministry of Education of the People's Republic of China. (2024). Guidance on Enhancing AI Literacy in Primary and Secondary Schools. http://en.moe.gov.cn/news/press_releases/202412/t20241210_1166454.html.
- Ministry of Education of the People's Republic of China. (2025). Guide to using Generative AI in Primary and Secondary Schools. (English translation by CSET) <https://cset.georgetown.edu/publication/china-school-gen-ai-use-guide-2025/>.
- National Institute of Standards and Technology (NIST). (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>.
- Ministry of Education of the People's Republic of China, et. al. (2026). AI + Education Action Plan. https://english.www.gov.cn/news/202604/15/content_WS69df29e6c6d00ca5f9a0a6b1.html.
- National Institute of Standards and Technology (NIST). (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>.
- Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. (2021). Conceptualizing AI Literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041 <https://doi.org/10.1016/j.caeai.2021.100041>.
- OECD. (2026). Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education https://www.oecd.org/en/publications/oecd-digital-education-outlook-2026_062a7394-en.html.
- Panagopoulou, F., Parpoula, C. & Karpouzis, K. (2025). Legal Perspectives on AI and the right to digital literacy in education. *Frontiers in Computer Science*, 7, 1692268

<https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2025.1692268/full>.

Park, J. -A., Kim, C. H., & Kim, H. -J., (2026). AI literacy for safe deployment: Cross-national evidence on the interaction between talent and governance. *Telecommunication Policy*, 50, 103106

<https://www.sciencedirect.com/science/article/pii/S0308596125002034>.

Pesek, I. (2026). Rethinking Education Governance in the Age of AI. *Frontiers of Digital Education*, 3(1), 11. <https://doi.org/10.1007/s44366-026-0085-z>.

Rizvi, F., & Lingard, B. (2010). *Globalizing Education Policy*. Routledge.

Selwyn, N., Hillman, T., Eynon, R., Ferreira, G., Knox, J., Macgilchrist, F., & Sancho-Gil, J. M. (2022). What's next for EdTech? Critical hopes and concerns for the 2020s. *Learning Media and Technology* 47(3), pp. 263-275
<https://doi.org/10.1080/17439884.2020.1694945>.

Swiecki, Z., Khosravi, H., Chen, G., Martinez-Maldonado, R., Lodge, J. M., Milligan, S., & Gasevic, D. (2022). Assessment in the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 3, 100075 <https://doi.org/10.1016/j.caeai.2022.100075>.

Taylor, S. (1997). Critical policy analysis: exploring contexts, texts and consequences. *Discourse: studies in the cultural politics of education*, 18(1), pp. 23-35
<https://doi.org/10.1080/0159630970180102>.

UNESCO. (2023). Guidance for Generative AI in education and research

<https://unesdoc.unesco.org/ark:/48223/pf0000386693>.

UNESCO. (2024). AI competency framework for students.

<https://www.unesco.org/en/articles/ai-competency-framework-students>.

UNESCO. (2024). AI competency framework for teachers

<https://www.unesco.org/en/articles/ai-competency-framework-teachers>.

US Department of Education, Office of Educational Technology. (2023). Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations
<https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf>.

US Department of Education. (2025). Guidance on the Use of Federal Grant Funds to Improve Education Outcomes Using Artificial Intelligence



<https://www.ed.gov/media/document/oepd-ai-dear-colleague-letter-7222025-110427.pdf>

US Department of Education. (2026). Final Priority and Definitions: Advancing Artificial Intelligence in Education. Federal Register. <https://www.federalregister.gov/documents/2026/04/13/2026-07087/final-priority-and-definitions-secretarys-supplemental-priority-and-definitions-on-advancing>.

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

White House. (2025). Executive Order 14277: Advancing Artificial Intelligence Education for American Youth <https://www.whitehouse.gov/presidential-actions/2025/04/advancing-artificial-intelligence-education-for-american-youth>.

White House. (2026). Remarks by Director Michael Kratsios at the India AI Impact Summit. February 20, 2026. <https://www.whitehouse.gov/releases/2026/02/remarks-by-director-michael-kratsios-at-the-india-ai-impact-summit/>.

Williamson, B., & Hogan, A. (2021). Commercialisation and privatisation in/of public education in the context of Covid-19. Education International https://www.researchgate.net/publication/343510376_Commercialisation_and_privatisation_in_of_education_in_the_context_of_Covid-19.

Yang, H., Xu, J., Zeng., & Gu, X. (2026). Comparing AI literacy policies in the European Union, the United States, India, and China. Telecommunications Policy, 50(3), 103142 <https://doi.org/10.1016/j.telpol.2025.103142>.