2026

Blue Book on Global K-12 AI Education Infrastructure

From Enabling Infrastructure to AI-Native K-12 Education: Literacy, Curriculum, Human–AI Collaboration, Assessment, and Governance

AI-SLI

2026 · AI-Assisted Research Edition

A Note on Method: How This Report Was Produced with AI Assistance

AI-SLI · Produced with AI Assistance

This blue book represents a systematic exploration by AI-SLI into AI-assisted knowledge production. Throughout its development, generative artificial intelligence served as a research assistant under the direction and oversight of the research team, undertaking four kinds of labour-intensive work. First, evidence-based research and data verification: fanning out multi-source searches across primary policy texts, international-organisation reports, authoritative surveys, and academic literature; tracing every policy-document number, survey statistic, and indicator cited in this volume to its primary source, cross-checking it, and issuing corrections where warranted; and systematically mapping the latest K-12 AI-education policies of China, the United States, the European Union, and international bodies such as UNESCO and the OECD. Second, theoretical anchoring and literature synthesis: mapping the domestic and international research lineage on literacy frameworks, curriculum design, human–AI co-learning pedagogy, process-oriented assessment, and governance mechanisms; comparing and cross-verifying differing frameworks and survey conventions across sources; and building on this basis the report's “literacy–curriculum–human–AI collaboration–assessment–governance” analytical framework. Third, drafting, figure generation, and citation management: composing the chapters and rendering the policy-comparison and data figures to a unified specification, and maintaining a fully traceable citation apparatus with graded credibility annotations. Fourth, a parallel bilingual edition: producing semantically aligned Chinese and English versions under a controlled terminology glossary.

We state plainly that the role of AI here was to carry out the labour-intensive work of searching, drafting, illustrating, and managing citations; the choice of subject, the judgements of value, the scholarly assessments, and the final conclusions were directed and vouched for by the research team. Every datum and policy cited is required to be real and independently verifiable, and every survey figure is required to state its issuing institution and methodology and must not be added across surveys with differing conventions. We offer this report as a forward-looking reference workflow for colleagues across education to scrutinise and improve upon — a sincere experiment in a new paradigm of knowledge production, and in no way a substitute for expert judgement or peer review.

Table of Contents

Chapter 1 Introduction: The Shift from "Enabling Infrastructure" to an AI-Native Paradigm for K-12 Education

Fig. 1 Teacher and student AI use has nearly doubled in two years

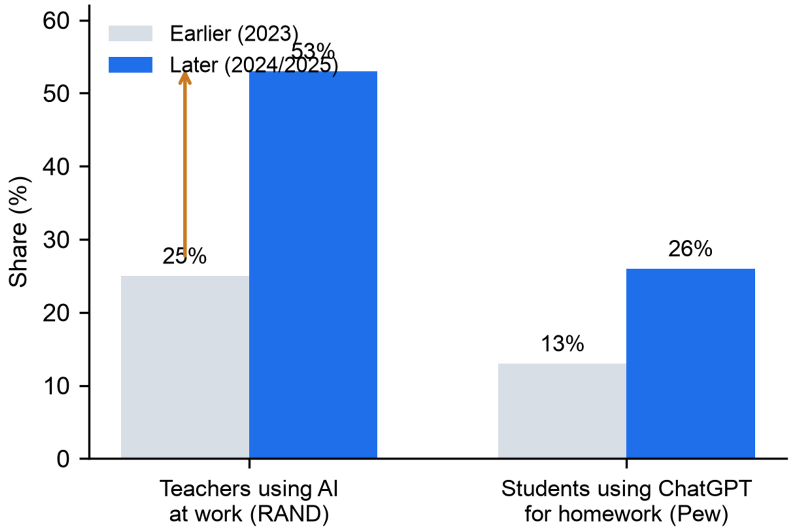

Source: RAND Corporation, American Educator Panels (2023-2025); Pew Research Center (2023-2024).

1.1 Origins: A Decade of Change Since the First White Paper

In 2022, this institute's predecessor published the *Global White Paper on Enabling Infrastructure for K-12 Artificial Intelligence Education*, organized around a "resources—teaching materials—algorithms—labs" throughline that systematically mapped the physical and content-level foundations that K-12 systems worldwide relied on to deliver AI education at the time. That edition was a product of a particular technological moment: large-scale generative AI had not yet entered public view, agents were still largely a laboratory concept, and "artificial intelligence" in the classroom mostly took the form of fixed algorithms, packaged tools, and purpose-built lab equipment. Back then, "enabling infrastructure" was chiefly about **giving teachers and students usable AI learning resources and practice venues** — open educational resources (OER), teacher communities, K-12 AI textbooks, algorithms and datasets, and AI labs, which together formed the six pillars of that earlier edition.

Several years on, the technological bedrock has shifted structurally, and the timeline of that shift is easy to date. ChatGPT's launch in November 2022 marked generative large language models' move from the lab into public life; by around 2023, education systems and international bodies everywhere found themselves forced to respond institutionally within an unusually short window. Generative models turned natural language into a universal interface; multimodal capability let text, image, speech, and code flow through a single model; on-device deployment began pushing inference power out of the cloud and onto personal devices and classroom terminals; and agentic AI turned the technology from a mere "answering tool" into a collaborator capable of planning, invoking other tools, and carrying out multi-step tasks. Together, these three forces — **agentic capability, multimodality, and on-device deployment** — have rewritten what "artificial intelligence" actually means inside a K-12 classroom.

This is not an abstraction — it is already visible in how teachers and students behave. In the United States, RAND Corporation's American Educator Panels survey found that the share of K-12 teachers using generative AI for work "doubled" across two school years, from roughly 25% in 2023–2024 to roughly 53% in 2024–2025[1]. Pew Research Center documented a parallel leap on the student side: among U.S. teens aged 13–17, the share who had used ChatGPT for schoolwork rose from 13% in 2023 to 26% in 2024 — doubling within a single year — with usage notably higher among older students (grades 11–12, at 31%) than younger ones (grades 7–8, at 20%)[2]. Once students can summon, on demand, an assistant that can converse, generate questions, grade work, and coach practice — and once teachers' own workflows for lesson planning, question-setting, and feedback are being reshaped by generative tools — the old "materials-supply" framework for enabling infrastructure simply no longer captures today's reality.

That is precisely why this institute undertook a systematic remake of the 2022 original, producing the *Global Blue Book on Enabling Infrastructure for K-12 Artificial Intelligence Education 2026*. This is not a data refresh. It is **a shift in analytical dimension**.

1.2 Core Thesis: Reconstructing What "Enabling Infrastructure" Means

The central claim of this Blue Book is this: in the generative and agentic era, the meaning of "enabling infrastructure" for K-12 AI education has moved from the **materials dimension** to the **capability–governance dimension**.

- **The old framework (2022, materials dimension)**: resources, textbooks, algorithms, datasets, labs — answering the question of "what to teach with, and where to teach it."
- **The new framework (2026, capability–governance dimension)**: literacy, curriculum, human–AI collaboration, assessment, governance — answering the question of "what capabilities to build, how to organize teaching, how humans and AI should work together, how to assess learning, and what rules should apply."

The deeper driver behind this reconstruction is straightforward: once AI tools become broadly accessible and their marginal cost of access approaches zero, the bottleneck constraining education quality is no longer "do we have the resources," but rather "do we have the literacy to work with AI well, the curriculum and assessment to match, and the governance to safeguard equity and integrity."

The critical elements of enabling infrastructure have therefore shifted — from **hard conditions supplied externally** to **soft capabilities and institutional arrangements embedded throughout the teaching process**.

This judgment rests on an observable empirical anchor: the central tension facing education systems today is not "insufficient adoption" but rather **adoption outpacing governance**. RAND's own surveys captured this succinctly under the heading "AI Use in Schools Is Quickly Increasing but Guidance Lags Behind": as of spring 2025, while roughly 45% of principals reported their school or district already had an AI use policy, only about 34% of teachers reported the existence of AI policies tied to academic integrity, and more than 80% of students said their teachers had never explicitly taught them how to use AI in their schoolwork[1][3]. In other words, the tools have already "trickled down" into the classroom, while the institutions meant to govern them are still catching up. That is the practical reason "enabling infrastructure" must move up from the materials dimension to the governance dimension — what is genuinely scarce today is no longer models or compute, but capability frameworks, curriculum design, assessment methods, and rules of use.

The table below lays out, side by side, the logic behind this shift between the two editions:

Dimension	**2022 original (materials dimension)**	**2026 Blue Book (capability–governance dimension)**	**Driver of the shift**
Starting point	Supply of resources and venues	Building capability and securing governance	Tool accessibility rose, moving the bottleneck upstream
Students	Using pre-built tools	AI literacy, prompt engineering, human–AI collaboration	Generative interaction became a universal interface
Teachers	Accessing textbooks and resources	Teacher AI literacy, collaboration with agentic teaching assistants	Lesson prep, question-setting, and feedback workflows were reshaped
Content	K-12 AI textbooks	Curriculum and syllabi mapped against literacy frameworks	From "textbook" to "curriculum system"
Assessment	[TBD: not yet systematized in the earlier edition]	AI literacy assessment, process-oriented assessment, academic integrity	Both the object and the method of assessment changed
Rules	Technical choices around data and algorithms	Data privacy, equity of access, use governance	Risk moved from the technical layer to the institutional layer

1.3 Three Independent Strands of External Evidence for the Paradigm Shift

"From materials to capability–governance" is not merely this institute's own judgment — it is corroborated by three independent strands of external evidence: a shift in international frameworks, a shift in major economies' policy, and a shift in market and classroom behavior. All three strands converge tightly within the 2023–2026 window, forming the evidentiary basis for this Blue Book's remake.

1.3.1 International Organizations: From "Cautious Warning" to "Competency Frameworks"

UNESCO's moves have been the most emblematic. In September 2023, UNESCO published its *Guidance for Generative AI in Education and Research* — the first global guidance document on generative AI specifically for education — advocating a "human-centred" approach that emphasized data privacy, ethical norms, and teacher capacity-building, while urging countries to set age thresholds and privacy protections for GenAI use in schools[4]. The tone at this stage still leaned toward "cautious warning."

The real turning point came in 2024. During UNESCO's 2024 Digital Learning Week, the organization released two competency frameworks in succession: the *AI Competency Framework for Teachers* and the *AI Competency Framework for Students*[5][6]. The teacher-facing framework proposed a structure spanning five dimensions and 15 competencies in total, explicitly grounded in principles of protecting teachers' rights and strengthening human agency and sustainability[5]. The arrival of these frameworks signaled that international discourse had moved on from "whether to use AI and how to guard against its risks" to "which competencies should be built, and how they should be woven into the curriculum" — which is exactly the core proposition of the capability–governance dimension.

Official documents from the United States reflect the same shift. As early as May 2023, the U.S. Department of Education's Office of Educational Technology published *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations*. Produced jointly with the nonprofit Digital Promise and drawing on four public hearings and input from more than 700 education stakeholders, the report staked out core positions such as "keeping humans in the loop," and offered recommendations including designing AI around modern learning principles, prioritizing trust-building, informing and involving educators, and developing education-specific guardrails[7]. This report was no longer answering "what resources to provide" — it was answering "what principles and rules should govern human–AI collaboration."

1.3.2 Major Economies: From "Scattered Pilots" to "Institutionalized Curriculum and Governance"

China, the United States, and the European Union all moved AI education into an institutionalized phase during 2024–2026, independently of one another — but by strikingly different paths, which is exactly what makes them the cross-national comparative backbone of this Blue Book.

China: advancing on twin tracks of a national curriculum guideline and use governance. On May 13, 2025, the Ministry of Education's Committee for Guiding Basic Education Teaching released the *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)*, built around a spiral curriculum design: at the primary level, the focus is experience and interest — "sensing the value of technology, understanding basic AI technologies such as speech recognition and image classification"; at the middle-school level, understanding and application are given equal weight, with students expected to grasp "the basic workflow of machine learning and the concept of supervised learning" and to complete data organization and agent-building tasks through project-based learning; at the high-school level, strategy and innovation are deepened further, requiring students to "understand the technical characteristics and social impact of generative AI," build algorithmic models, and develop integrated solutions[8]. Issued alongside it was the *Guidelines for Generative AI Use by Primary and Secondary Students (2025 Edition)*, together forming a dual-track policy structure of "universal literacy education plus use governance"[9]. Then, on April 2, 2026, five ministries and commissions — the Ministry of Education, the National Development and Reform Commission, the Ministry of Industry and Information Technology, the Ministry of Science and Technology, and the National Data Administration — jointly issued the *Action Plan for "AI + Education"* (Document No. 1 [2026] of the Ministry of Education's Department of Science, Technology and Informatization; Chinese: 教科信〔2026〕1 号), which further calls for "continuously refining the Guidelines for General AI Education in Primary and Secondary Schools, and ensuring AI-related courses are fully and properly offered," directs localities to "develop AI curriculum guidelines that specify course objectives, content, and class-hour requirements for each school stage," and stresses strengthening K-12 AI education bases and supporting schools in rural and remote areas in using national platforms to deliver these courses well[10]. What distinguishes China's approach is folding AI education into the national curriculum system while simultaneously laying down use norms — "capability" and "governance" advanced together, in a single stroke, through central-government documents.

The United States: federal guidance as soft steer, with states and districts as the point of implementation. Unlike China's top-down curriculum guidelines, the U.S. has no unified national curriculum at the K-12 level; the federal Department of Education's role is largely limited to publishing insights, recommendations, and "guardrail" frameworks[7], while the concrete policies, training, and academic-integrity rules are left to states and districts. RAND's surveys paint a clear picture of the resulting tension within this decentralized system: on one hand, adoption is climbing fast (teacher work-related usage rose from roughly 25% to roughly 53% within two years)[1]; on the other, institutional support and training are lagging conspicuously and unevenly distributed — the share of districts offering teachers AI training rose from 23% in fall 2023 to 48% in fall 2024, yet low-poverty districts (about 67%) far outpaced high-poverty districts (about 39%)[1][3]. The American path thus exhibits a distinctive pattern of "behavior first, governance playing catch-up, and uneven regional distribution," with equity and academic integrity emerging as the most prominent governance concerns.

The European Union: hard legal constraints, with education classified as a high-risk domain. The EU has taken the path of "framing through regulation." The EU AI Act took effect on August 1, 2024, and explicitly places "education and vocational training" on its list of high-risk applications — any AI system used to determine access to or admission into education, to assess learning outcomes, to determine the appropriate level of education for an individual, or to monitor and detect prohibited behavior during tests is classified as high-risk and must satisfy strict compliance requirements covering transparency, accuracy, human oversight, and risk management[11]. The Act applies in stages: prohibited AI practices and AI literacy obligations took effect on February 2, 2025, while the rules governing high-risk domains such as education apply from 2027 (per the EU's own timeline for delayed application of the high-risk provisions)[11]. What distinguishes the EU's approach is that it front-loads "governance" as a legally binding compliance framework, while simultaneously writing "capability" requirements into law through its AI literacy obligations.

Read together, these three cases show that regardless of whether the lever is a curriculum guideline (China), federal soft guidance layered with state/district implementation (the U.S.), or hard legal constraint (the EU), the destination is the same: away from "providing tools" and toward "building capability plus establishing rules." This gives the Blue Book's "capability–governance" framework consistent evidentiary support across very different institutional systems.

The table below summarizes how the three systems' paths diverge:

System	Primary lever	Representative documents (2023–2026)	Positioning of curriculum/capability	Governance emphasis
China	National curriculum guideline + dual-track use norms	*Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)*; *Guidelines for Generative AI Use by Primary and Secondary Students	Stage-by-stage spiral curriculum on general AI literacy, folded into the national curriculum system	Universal literacy education and use norms advance together, with urban–rural balance in view

		(2025 Edition)*; *Action Plan for "AI + Education"* (Document No. 1 [2026])[8][9][10]		
United States	Federal soft guidance + state/district implementation	Department of Education, *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations* (2023)[7]	No unified national curriculum; capability requirements steered by frameworks, implemented by districts	Academic integrity, training supply, and regional equity
European Union	Hard legal constraint	EU AI Act (took effect 2024)[11]	Capability requirements carried through statutory AI literacy obligations	Education classified as high-risk, with mandatory compliance
International organizations	Global framework supply	UNESCO *Guidance for Generative AI in Education and Research* (2023); *AI Competency Framework* for teachers and students (2024)[4][5][6]	15 teacher competencies and dimension-based student competency lists	Human-centeredness, privacy, agency

1.3.3 Market and Classroom Behavior: From "Purpose-Built Tools" to "General-Purpose Teaching Assistant"

The third strand of evidence comes from the market and from what is actually happening in classrooms — and it explains why governance pressure has concentrated so sharply between 2024 and 2026.

First, leading AI vendors have made "education" a standalone product line, and one that is spilling over rapidly from higher education. In 2024, OpenAI launched ChatGPT Edu for universities; in April 2025, Anthropic released Claude for Education, which includes a built-in "Learning Mode" that guides students toward reasoning through prompted questions rather than handing them direct answers, and struck campus-wide partnerships with universities including Northeastern University — a single deployment that alone covered roughly 50,000 students, faculty, and staff[12][13]. Around the same period, the industry saw a collective pivot toward "learning/guided" modes — ChatGPT's Study Mode and Google Gemini's Guided Learning both launched[12]. This evolution in product design is itself a governance signal: vendors have begun proactively drawing pedagogical boundaries around whether to hand over direct answers at all.

Second, students and teachers alike have effectively turned general-purpose chat assistants into de facto "teaching assistants." Pew's data show student-side usage doubling within a year (13% to 26%)[2]; RAND's data show teacher-side usage doubling within two years (roughly 25% to 53%)[1]. Notably, most teachers are not using education-specific platforms at all, but general-purpose ChatGPT[3] — meaning that the AI actually running in classrooms is often not purpose-built for K-12 settings and may lack built-in age and privacy guardrails. This is precisely the real-world reason UNESCO and the EU keep returning to privacy, age thresholds, and compliance[4][11].

Third, students show a clear sense of where "acceptable use" ends, which offers a behavioral foundation for assessment and academic-integrity governance. Pew's survey found teens far more comfortable using ChatGPT "to learn something new/for research" (54% saying this is acceptable) than for "solving math problems" (29%) — and considerably less comfortable still with using it "to write an essay" (18%)[2]. This pattern shows that students are not indiscriminately "cheating with AI"; rather, their comfort with use declines along a spectrum from "aiding understanding" to "substituting for original output" — precisely the zone where process-oriented assessment and integrity governance can gain real traction.

Taken together, these three strands of evidence converge within the same window of time, jointly supporting this Blue Book's core judgment: the critical elements of enabling infrastructure have moved up from "materials" to "capability–governance."

1.4 The Three Technological Undercurrents of 2026

The paradigm shift described above is happening now because of the technological reality of 2026. Every argument this Blue Book makes about enabling infrastructure rests on a basic reading of three technological undercurrents. What follows is a qualitative sketch only; the specific scale of adoption and the product landscape are addressed in later chapters and in this institute's companion Blue Books, cross-referenced accordingly.

1.4.1 Agentic AI

AI is shifting from passive respondent to active collaborator. Agents can now understand instructional goals, break tasks down, invoke other tools, execute multi-step processes, and self-correct along the way — making "agentic teaching assistants" practically viable for answering questions, generating exercises, personalized coaching, and first-pass grading of assignments. The industry-wide pivot toward "learning/guided" modes is an early form of agentic AI being constrained pedagogically for education: its value orientation is shifting from "completing the work for the student" to "accompanying the student's thinking"[12][13]. This brings real potential to reduce workload and improve efficiency, but it also raises new governance questions about teachers' role, the boundary between student and AI, and academic integrity — with nearly half of students using AI on their own without clear guidance[1], "how to collaborate" and "where to draw the line" become curriculum and governance questions at once. For a related discussion of first-person classroom interaction, see this institute's *Blue Book on AI Glasses in Education 2026*.

1.4.2 Multimodality

Text, image, speech, video, and code now flow through a single unified model, so AI-enabled support is no longer confined to text-based Q&A. Voice interaction for younger learners, accessibility support for students with special needs, and visual understanding for lab and hands-on courses have all gained new points of entry thanks to multimodal capability. The U.S. Department of Education's 2023 report already anticipated this, noting that AI's ability to communicate through speech, gesture, and sketching could open new forms of support for students with disabilities, for collaborative learning, and for managing complex classroom routines[7]. This undercurrent matters especially for equity and accessibility, discussed further in the chapter on equity of access.

1.4.3 On-Device Deployment

Inference capability is moving down toward personal devices, classroom terminals, and local servers, bringing lower latency, stronger privacy protection, and usability in low- or no-connectivity environments. On-device deployment bears directly on how the urban–rural digital divide can be closed and on how student data can be governed locally — on this point, China's *Action Plan for "AI + Education"* specifically calls for "supporting schools in rural and remote areas in using national platforms to deliver relevant courses well"[10], while RAND's data on the training gap between high- and low-poverty districts (roughly 67% vs. 39%) illustrates, from the opposite direction, just how urgent equity-focused governance has become[3]. On-device deployment is the technical fulcrum underpinning both the "accessible" and the "trustworthy" dimensions of enabling infrastructure. For educational robots as the physical, on-device embodiment of agentic AI, see this institute's *Global White Paper on Educational Robot Development 2026*.

> Note: More granular quantitative indicators for these three undercurrents — market size, unit shipments, model parameters, and region-by-region adoption rates — are presented with full sourcing in the corresponding later chapters. This chapter cites only the authoritative qualitative judgments and the small set of verifiable data points directly relevant to the paradigm shift itself.

1.5 Scope and Boundaries of This Study

- **Grade-level scope**: This Blue Book focuses on basic education (K-12 / primary and secondary schooling), with necessary notes on the pre-school transition and the handoff into higher education where relevant, though neither is treated as a primary focus. It should be noted that some of the first-hand adoption data cited (such as campus-wide deployments of Claude for Education and ChatGPT Edu) originated in higher-education settings; this chapter cites such data only as evidence for the broader trend of "general-purpose teaching assistants becoming productized," and this does not shift the Blue Book's primary K-12 focus[12][13].

· **Geographic scope**: **China, the United States, and the European Union** serve as the primary cross-national comparative axis for policy and curriculum systems (see the comparison table in Section 1.3.2 for how their paths diverge), supplemented by representative practices from other notable countries and regions. Fuller comparisons of policy document numbers, curriculum names, and release dates appear in the later chapters on curriculum and syllabi.

· **Thematic scope**: This Blue Book covers AI literacy, curriculum and syllabi, human–AI collaborative teaching, assessment and academic integrity, equity of access, teacher AI literacy, and data and use governance — the full set of elements within the capability–governance dimension.

· **Explicit exclusions**: This Blue Book does not cover AI talent development as a higher-education specialty, corporate training, or purely academic research unrelated to the K-12 mainline; nor does it offer selection recommendations for specific commercial products (any mention of specific products in this chapter serves only as evidence of a trend, not as an endorsement).

1.6 Methodology and Evidentiary Principles

This Blue Book adheres to the principles of **research first, evidence-based visualization, and traceable sourcing**:

1. **Framework first**: The analytical framework of "literacy → curriculum → human–AI collaboration → assessment → governance" is established first, and evidence and narrative are then organized around it.
2. **Cross-national comparison**: Policy, curricula, and practice in China, the United States, and the EU are compared side by side in a structured way, avoiding the trap of letting any single system stand in for the world (Section 1.3.2 of this chapter demonstrates this comparative method).
3. **Evidence-based visualization**: Available data on adoption, investment, and coverage are presented visually to show trends and differences; wherever a specific figure, share, document number, or citation is used, its real source is noted, and any gap is marked explicitly as [TBD: data/source] rather than invented. All data cited in this chapter is drawn from primary sources (Pew, RAND, UNESCO, China's Ministry of Education, the EU, the U.S. Department of Education, and others) and cross-checked wherever possible.

4. **Temporal updating**: Data from the 2024–2026 window on policy and adoption replaces the pandemic-era data used in the 2022 original, ensuring currency.
5. **Careful attention to data definitions**: Because different surveys vary in sample, grade level, and definition of "use" (personal use vs. instructional use; any use vs. frequent use), their figures cannot simply be added together or treated as equivalent. This Blue Book notes the source and definitional scope wherever it cites data, to avoid any "additive exaggeration."

Data and sourcing convention: All specific statistical definitions, sample scopes, and citations used throughout this Blue Book are consolidated in each chapter's "Chapter References" section and in the book's overall bibliography; see [TBD: bibliography and data appendix].

1.7 A Guide to the Structure of This Blue Book

Following this chapter, the Blue Book unfolds along the "capability–governance" throughline (specific chapter titles and numbers are subject to the final table of contents):

Chapter	**Topic**	**Core question**
Chapter 1	Introduction and the paradigm shift	Why the shift from the materials dimension to the capability–governance dimension
Chapter [TBD]	AI literacy framework	Which AI competencies K-12 students should possess
Chapter [TBD]	Curriculum and syllabi compared	How China, the U.S., and the EU organize AI curricula
Chapter [TBD]	Human–AI collaboration and agentic teaching assistants	How teachers and students collaborate with AI in instruction
Chapter [TBD]	Assessment and academic integrity	How to assess AI literacy and safeguard integrity
Chapter [TBD]	Equity of access	How to close the urban–rural divide and support special populations
Chapter [TBD]	Teacher AI literacy	What AI capabilities and support teachers need
Chapter [TBD]	Data and use governance	What rules should govern AI use on campus
Chapter [TBD]	Conclusions and outlook	Where enabling infrastructure goes next

1.8 Chapter Summary

From 2022 to 2026, what has changed is not just the technology, but our very way of understanding "enabling infrastructure." Three independent strands of evidence converge within the same window of time: international organizations moved from "cautious warning" to "competency frameworks" (UNESCO's 2023 guidance, followed by its 2024 teacher and student competency frameworks)[4][5][6]; major economies moved from "scattered pilots" to "institutionalization" (China's national curriculum guideline paired with dual-track governance, the U.S.'s federal soft guidance layered with state/district implementation, and the EU's legal classification of education as high-risk)[7][8][9][10][11]; and market and classroom behavior moved from "purpose-built tools" to "general-purpose teaching assistant," with adoption doubling within a year or two while governance and training clearly lagged behind and remained unevenly distributed[1][2][3][12][13].

The picture this paints points in one clear direction: once artificial intelligence shifts from being a classroom **tool** to being a classroom **collaborator**, what is genuinely scarce is no longer the model, the compute, or the resources — it is **the literacy to work with AI well, the curriculum to organize teaching around it, and the governance to safeguard equity and integrity**. The critical work of supporting education has therefore shifted from "what resources to provide" to "what capabilities to build and what rules to establish." Organized around the throughline of "literacy → curriculum → human–AI collaboration → assessment → governance," this Blue Book aims to offer a comprehensive picture of K-12 AI education in the generative and agentic era — one that holds capability-building and governance safeguards in equal view. The chapters that follow will develop this throughline one by one, answering the core questions listed in the table above through evidence-based, traceable, cross-nationally compared analysis.

Chapter References

[1] AI Use in Schools Is Quickly Increasing but Guidance Lags Behind: Findings from the RAND Survey Panels · RAND Corporation · 2025 · https://www.rand.org/pubs/research_reports/RRA4180-1.html

[2] About a quarter of U.S. teens have used ChatGPT for schoolwork – double the share in 2023 (Pew survey, September 18–October 10, 2024, sample of 1,391 teens aged 13–17) · Pew Research Center · 2025 · https://www.pewresearch.org/short-reads/2025/01/15/about-a-quarter-of-us-teens-have-used-chatgpt-for-schoolwork-double-the-share-in-2023/

[3] Uneven Adoption of Artificial Intelligence Tools Among U.S. Teachers and Principals in the 2023–2024 School Year (American Educator Panels, Research Report RR-A134-25) · RAND Corporation · 2025 · https://www.rand.org/pubs/research_reports/RRA134-25.html

[4] Guidance for Generative AI in Education and Research · UNESCO · 2023 · https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

[5] AI Competency Framework for Teachers · UNESCO · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-teachers

[6] AI Competency Framework for Students · UNESCO · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-students

[7] Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations · U.S. Department of Education, Office of Educational Technology · 2023 · https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf

[8] Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [中小学人工智能通识教育指南（2025 年版）] (issued by the Ministry of Education's Committee for Guiding Basic Education Teaching, May 13, 2025) · China Education and Research Network (CERNET) · 2025 · https://www.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[9] Ministry of Education's Committee for Guiding Basic Education Teaching Releases the Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) and the Guidelines for Generative AI Use by Primary and Secondary Students (2025 Edition) [教育部基础教育教学指导委员会发布《中小学人工智能通识教育指南（2025 年版）》和《中小学生成式人工智能使用指南（2025 年版）》] · China Education Online and related reporting · 2025 · https://www.canedu.org.cn/site/content/7876.html

[10] Notice of the Ministry of Education and Four Other Departments on Issuing the Action Plan for "AI + Education" [教育部等五部门关于印发《"人工智能+教育"行动计划》的通知] (Document No. 1 [2026]; Chinese: 教科信〔2026〕1 号, April 2, 2026) · Ministry of Education of the People's Republic of China · 2026 · http://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html

[11] Regulatory framework on Artificial Intelligence (EU AI Act, took effect August 1, 2024; education classified as high-risk; AI literacy obligations apply from February 2, 2025) · European Commission · 2024 · https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[12] OpenAI, Anthropic target college students with latest education AI announcements (covers the ChatGPT Edu, Claude for Education, and Study Mode / Guided Learning timeline) · CNBC · 2025 · https://www.cnbc.com/2025/04/03/openai-anthropic-target-college-students-with-education-ai-services.html

[13] Introducing Claude for Education (covers Learning Mode and the roughly 50,000-person campus-wide deployment at Northeastern University) · Anthropic · 2025 · https://www.anthropic.com/news/introducing-claude-for-education

Chapter 2 The Global Policy Landscape: Comparing K-12 AI Education Regulations and Curriculum Standards in China, the United States, and the EU (2024–2026)

Fig. 2 K-12 AI education policy timeline: China, the United States, the EU and UNESCO (2023-2026)

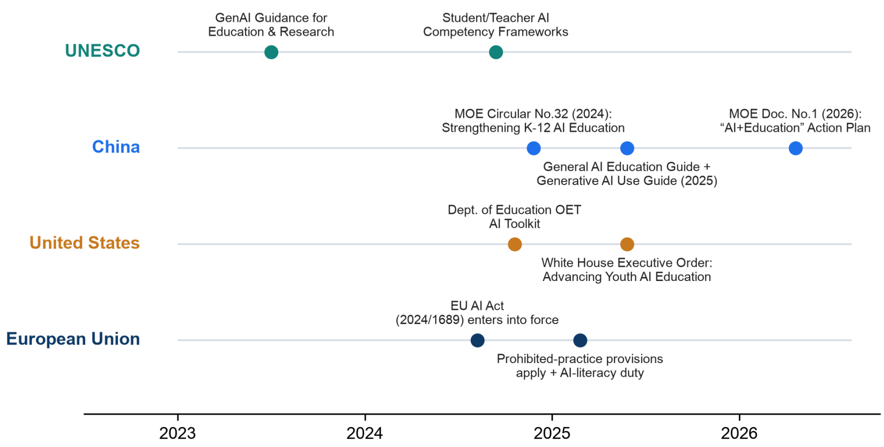

Source: China's Ministry of Education; US Dept. of Education OET and the White House; EU AI Act (2024/1689); UNESCO (see Chapter References).

2.1 Introduction: From "Informatization Policy" to "AI Education Governance"

The 2022 master edition of this Blue Book was drafted before generative AI entered the public consciousness at scale. At that time, most countries' K-12 policies treated AI as an extension of existing "educational informatization," "digital literacy," and "STEM/computational thinking" frameworks — a frontier topic schools might optionally teach, or a resource reserved for the lab. Since late 2022, the rapid spread of generative AI, multimodal large models, and agent-based systems has reshaped that policy context in a structural way: AI is no longer simply a curricular question of whether to teach it, but has become simultaneously a governance question of how to use it safely, equitably, and responsibly on campus.

That shift shows up across China, the United States, and the EU alike between 2024 and 2026, all three moving their center of gravity from **supplying content** (what to teach, what resources to use) toward **balancing capability-building with governance** (what competencies to cultivate, how to regulate use, how to assign accountability). Each jurisdiction's flagship documents bear this out. China's Ministry of Education issued the *Notice on Strengthening Artificial Intelligence Education in Primary and Secondary Schools* (Document No. 32 [2024] of the General Office of the Ministry of Education; Chinese: 教基厅函〔2024〕32 号) in November 2024, setting a quantified target of making AI education essentially universal in primary and secondary schools by 2030 [1][2]. The United States moved first through the Department of Education's Office of Educational Technology, which released an AI-integration toolkit for district leaders in October 2024, followed by a White House executive order, *Advancing Artificial Intelligence Education for American Youth*, signed in April 2025 [5][6][9]. The EU, meanwhile, anchored its approach in the Artificial Intelligence Act (Regulation (EU) 2024/1689) — adopted in June 2024 and in force by August — paired with the *Digital Education Action Plan 2021–2027* and its literacy frameworks, forming twin pillars of "regulation plus literacy" [11][14]. All three timestamps fall after the generative-AI shock wave, marking a transition for AI education governance from the "informatization materials" phase into a phase of institution-building.

Rather than attempt an exhaustive catalogue of every national document, this chapter builds a comparable analytical framework along four threads — legislative and policy tier, curriculum standards, teacher preparation and implementation, and governance and compliance — to compare the three jurisdictions side by side, identify convergences, divergences, and institutional gaps, and use UNESCO's global frameworks as a supranational reference point. Every specific document number, publication date, grade-band coverage, and quantitative figure cited in this chapter is grounded in verified primary or authoritative sources; anything not yet verified is marked with a `[TBD: ...]` placeholder rather than invented.

> This chapter cross-references other chapters of AI-SLI's *Blue Book on the Global Enabling Infrastructure for K-12 Artificial Intelligence Education 2026*: competency-tier detail for literacy goals appears in Chapter 3 (Curriculum and Literacy Frameworks); implementation detail on teacher AI literacy appears in the teacher-focused chapter; academic integrity and assessment governance appear in the assessment and governance chapter. For policy implications at the level of specific device categories, see AI-SLI's *AI Glasses in Education White Paper 2026* and *Global Development of Educational Robots White Paper 2026*.

2.2 A Framework for Comparison

To make the three jurisdictions' policies comparable, this chapter sets out four comparative dimensions and a set of observation points beneath each.

2.2.1 Four comparative dimensions

- **Dimension One — Policy and legislative tier**: the legal standing of AI-education rules (statute, administrative regulation, ministerial rule, guidance, or industry guideline), and whether authority flows top-down from a unified national mandate or bottom-up from local/district autonomy.

- **Dimension Two — Curriculum and literacy standards**: whether a jurisdiction has a standalone K-12 AI curriculum standard or instead embeds AI literacy within existing subjects; how literacy goals are structured (knowledge–skills–attitudes–ethics) and how they progress across grade bands.
- **Dimension Three — Teacher preparation and implementation support**: requirements for teacher AI literacy, pre-service and in-service training mechanisms, instructional resources and infrastructure, and implementation pacing (pilot-then-scale).
- **Dimension Four — Governance, ethics, and compliance**: data protection and minors' privacy, boundaries on generative AI use in the classroom, academic integrity, algorithmic transparency and accountability, and safety/compliance review of AI systems used by students.

2.2.2 A qualitative sketch of the three governance paradigms

Without presupposing that any one approach is superior, the three jurisdictions' governance logics can be characterized as follows (supporting documents are detailed in the sections that follow):

- **China**: driven by unified national deployment, with education authorities issuing top-down guidance that folds AI education into the curriculum system and rolls it out progressively by grade band, dovetailing with the broader digital-education strategy, the "AI+ Education" initiative, and efforts to cultivate top innovative talent — all backed by explicit timetables and quantified targets [1][3].
- **United States**: characterized by federal tone-setting with implementation left to states and districts. Federal-level action (the Department of Education's toolkit, the White House executive order) is largely directional guidance and incentive-based rather than a binding curriculum mandate; substantive curriculum standards and generative-AI use policies are set independently by states and districts, producing a highly fragmented, varied landscape [5][6][7].
- **European Union**: built on twin pillars of horizontal legislation (the AI Act) and digital-education action, using risk-tiered regulation and literacy frameworks to shape member states' compliance boundaries and literacy requirements from the top down — while leaving actual curriculum design to member-state sovereignty [11][14][17].
- **UNESCO (supranational reference)**: non-binding, but its *Guidance for Generative AI in Education and Research* (2023) and *AI Competency Frameworks for Students/Teachers* (2024) function as a global public good, supplying common normative language and competency benchmarks that profoundly shape policy discourse in all three jurisdictions — particularly in developing countries [19][21][22].

2.3 Policy Details by Jurisdiction

2.3.1 China: Top-down national deployment with a quantified timetable

(1) Top-level deployment and document number. The centerpiece of China's K-12 AI education policy is the *Notice on Strengthening Artificial Intelligence Education in Primary and Secondary Schools*, issued by the General Office of the Ministry of Education on November 21, 2024 (Document No. 32 [2024]; Chinese: 教基厅函〔2024〕32 号) [1][2]. This is the first national-level document to set an overarching goal — making AI education essentially universal in primary and secondary schools before 2030 — and it lays out six major tasks: building a systematic curriculum framework, implementing regular instruction and assessment, developing broadly applicable teaching resources, establishing pervasive learning environments, expanding the teacher pipeline at scale, and organizing diverse exchange activities [1][2]. This is a typical "guidance-based deployment" document — a normative document ranking below formal ministerial rules — whose force comes from administrative coordination rather than legislative constraint, though it carries substantial mobilizing power when paired with the central digital-education strategy.

(2) Curriculum positioning and grade-band progression. The Notice establishes a clear, progressive structure across grade bands [1][2]:

- Lower primary grades: emphasis on perceiving and experiencing AI technology;
- Upper primary grades and middle school: emphasis on understanding and applying AI technology;
- High school: emphasis on project-based creation and frontier applications.

On curriculum delivery, China follows a flexible, dual-track approach — "standalone courses running alongside subject integration" — relying primarily on Information Technology and General Technology courses while also allowing AI content to be woven into science, comprehensive practical activities, and labor education [1].

(3) Piloting first: 184 education bases. Ahead of the Notice, the Ministry of Education had already, under the *Notice on Recommending AI Education Bases in Primary and Secondary Schools* (Document No. 29 [2023]; Chinese: 教基厅函〔2023〕29 号), run a provincial-selection and public-announcement process that culminated in February 2024 with **184** designated AI education bases in primary and secondary schools — serving as pilot anchors for exploring new concepts and models and generating scalable experience [3]. This exemplifies China's characteristic "pilot-then-scale" policy rhythm.

(4) A local implementation case: Beijing. Beijing stands out as an early adopter of the national mandate. Building on Document No. 32 [2024] and the *Beijing Work Plan for Advancing AI

Education in Primary and Secondary Schools (2025–2027)*, the Beijing Municipal Education Commission issued the *Beijing Local Curriculum Outline for AI Education in Primary and Secondary Schools (Trial) (2025 Edition)* on June 26, 2025, stipulating that starting in **fall semester 2025**, all primary and secondary schools citywide will offer general AI education, with **no fewer than 8 class hours per academic year**, spanning every grade band from primary through high school [4][10]. The course can be offered standalone or blended with Information Technology, General Technology, or Science; grade-band emphasis runs from early exposure in primary school, to cognitive empowerment in middle school, to comprehensive practice and innovation in high school [4]. On the teacher-supply side, Beijing has launched the "Hundred-Thousand Seed Program," aiming to cultivate roughly 100 expert-level teachers and 1,000 backbone teachers [10].

(5) The 2026 strategic upgrade: the "AI+ Education" Action Plan. In April 2026, the Ministry of Education, jointly with the National Development and Reform Commission, the Ministry of Industry and Information Technology, the Ministry of Science and Technology, and the National Data Administration, issued the *"AI+ Education" Action Plan* (Document No. 1 [2026] of the Ministry of Education, Science and Technology Department; Chinese: 教科信〔2026〕1 号) [3]. Targeting 2030, the plan calls for "building a vertically integrated, horizontally connected AI education system spanning all grade levels and a general AI literacy system for society at large," and lays out four priority tasks: advancing AI talent cultivation and literacy development, deepening and broadening the integration of AI with education, strengthening the foundational environment for "AI+ Education," and optimizing the development ecosystem [3]. Notably, the plan places "accelerating universal AI education for primary and secondary students and folding it fully into local curriculum systems" alongside "incorporating AI into teacher qualification examinations and certification" — signaling that China has moved beyond a single-point curriculum requirement toward an integrated system spanning talent cultivation, teacher certification, foundational infrastructure, and industrial ecosystem [3].

(6) Governance orientation. While encouraging adoption, Chinese policy also emphasizes moral education and a human-centered approach, and sets normative requirements around generative AI entering schools and protecting student data and minors. As of this chapter's research cutoff, however, independent national-level rules specifically governing generative-AI use in the classroom remain largely at the level of guidance-style language rather than binding detail [TBD: document number/provisions of any dedicated national rule on generative AI entering schools].

2.3.2 United States: Federal tone-setting, fragmented state and district implementation

(1) Federal guidance: a non-mandatory directional framework. The federal government sets no mandatory curriculum standard for K-12 AI education. In October 2024, the Department of

Education's Office of Educational Technology (OET) released *Empowering Education Leaders: A Toolkit for Safe, Ethical, and Equitable AI Integration*, a roughly 74-page resource aimed at district leaders, organized around three principles — safety, ethics, and equity — with a dedicated module underscoring the importance of educator AI literacy [5][6]. Earlier, in May 2023, the same office's report *Artificial Intelligence and the Future of Teaching and Learning* had already established principles such as human-centered design and "keeping humans in the loop." All of these documents are guidance in nature and carry no legal force.

(2) The 2025 executive order: reinforcing federal incentives. On April 23, 2025, the President signed the executive order *Advancing Artificial Intelligence Education for American Youth*, establishing a White House AI Education Task Force chaired by the Director of the Office of Science and Technology Policy, with members drawn from the Departments of Education, Labor, Energy, and Agriculture, the National Science Foundation, and others [6][9]. The order launches a "Presidential AI Challenge," promotes public-private partnerships to develop online K-12 AI literacy resources, and directs the Secretary of Education to issue guidance, within a set period, on how grant funding can be used to improve AI-related education outcomes [6][9]. This further reinforces the federal government's governance logic of steering through incentives and coordination rather than mandating uniformity.

(3) State and district autonomy: fragmented but spreading fast. The real authority to set K-12 AI curricula and generative-AI use policy sits with states and districts. According to ongoing tracking by organizations such as AI for Education, the number of states that have issued official AI-education guidance or policy has climbed rapidly — multiple 2025 tallies place the figure between roughly **33 and 35 states** (AI for Education reported "31 states" in 2025 and later updated that to 33; the National Association of State Boards of Education (NASBE) counted roughly 34 states by the end of 2025; another tally that includes Puerto Rico reaches 35) [7][8]. Representative state-level actions include Georgia (January 2025), Maine (an interactive guidance toolkit released February 2025), New Mexico (Version 1.0 released May 2025), Massachusetts (August 2025), and Montana (October 2025) [7]. State guidance documents vary considerably in structure and binding force, with most focused on principles of privacy, equity, and responsible use rather than mandatory curriculum standards.

(4) Non-governmental sources of literacy frameworks. Because the federal government supplies no unified curriculum standard, a substantial share of U.S. K-12 AI literacy goals are drawn from frameworks proposed by non-governmental organizations and industry coalitions — such as the TeachAI coalition (launched by Code.org, ETS, ISTE, and others), which offers a reference template for educators in states lacking official guidance of their own [7][8]. The result is a distinctive supply structure: "loose federal guidance, varied state action, NGOs filling the gaps."

(5) Governance orientation: existing privacy law layered with state legislation. The U.S. has no dedicated federal AI-education law; governance of student data instead relies on existing federal statutes — the Family Educational Rights and Privacy Act (FERPA) protects student education records, while the Children's Online Privacy Protection Act (COPPA, enforced by the Federal Trade Commission) sets parental-consent requirements for collecting personal information from children under 13 [23]. On top of that baseline, roughly **40 states** have enacted their own student-data privacy laws, producing a patchwork landscape [23]. Compliance gaps are pronounced: a 2024 survey by the Center for Democracy & Technology (CDT) found that among districts using AI tools, 42% had not signed a data-processing agreement with their AI vendor, and 31% of administrators could not identify which federal law applied to student data privacy [23].

2.3.3 European Union: Layered governance through horizontal legislation plus digital-education action

(1) The horizontal regulatory framework: the Artificial Intelligence Act. The EU's overarching AI governance instrument is the Artificial Intelligence Act (Regulation (EU) 2024/1689). It was adopted on June 13, 2024, published in the EU's Official Journal on July 12, and **entered into force on August 1, 2024**, with staggered application dates: prohibitions and the "AI literacy obligation" apply from **February 2, 2025**; obligations for general-purpose AI (GPAI) models apply from August 2, 2025; obligations for the high-risk systems listed in Annex III were originally set to apply from August 2026, but a 2026 provisional agreement under the "Digital Omnibus" package pushed that deadline from August 2, 2026 to **December 2, 2027** [11][15].

(2) A dual positioning for education: both "prohibited" and "high-risk." This is the part of the EU framework most directly relevant to education — and also the most frequently misread — and the two categories must be kept strictly distinct [11][12]:

- **Prohibited (Article 5)**: AI systems used to **infer the emotions** of natural persons **in the workplace and in educational institutions** are classified as a "prohibited AI practice," permitted only under narrow exceptions such as medical or safety purposes; this prohibition took effect on February 2, 2025 [11]. In practice, this means classroom "emotion recognition/attention monitoring" applications are, in principle, banned across the EU.
- **High-risk (Annex III)**: Specific AI uses in education and vocational training — for example, determining admission or enrollment, evaluating learning outcomes, assessing the appropriate level of education, and monitoring for exam cheating — are classified as high-risk systems, subject to strict compliance obligations covering risk management, data governance, log retention, and human oversight [12].

In other words, within the same "education" domain, emotion inference is **prohibited** outright, while admissions- and assessment-related AI is **high-risk** and subject to heavy regulation — the two categories should never be conflated.

(3) Codifying the AI literacy obligation. Article 4 of the AI Act establishes an "AI literacy" obligation, applicable from February 2, 2025, requiring providers and deployers of AI systems (which, in an educational context, can include schools and educational institutions) to take measures ensuring that relevant personnel have a sufficient level of AI literacy [11]. This elevates "AI literacy" from a policy aspiration to a binding legal requirement — a step unique to the EU among the three jurisdictions.

(4) Digital-education action and literacy frameworks. Running alongside hard legislation, the EU uses the *Digital Education Action Plan 2021–2027* (adopted September 30, 2020) as its strategic framework for driving member states' digital transformation [14]. Two key literacy and ethics instruments sit beneath it:

- The **Ethical guidelines on the use of AI and data in teaching and learning for educators** (published October 2022), produced by the Commission's Expert Group on AI and Data in Education and Training (active July 2021–June 2022), aimed at primary and secondary school teachers to help them cultivate students' ethical awareness and critical thinking while complying with instruments such as the AI Act [17][18].
- **DigComp 2.2: The Digital Competence Framework for Citizens** (published by the Joint Research Centre (JRC) in March 2022), which retains the framework's five competence areas and 21 competencies while adding more than **250** new examples of knowledge, skills, and attitudes — a large share of which point directly at AI and automated systems, such as recognizing AI-generated content, understanding algorithmic recommendations, and recognizing AI bias [16]. DigComp has thus become the EU's shared reference point for embedding AI literacy within civic digital-competence frameworks.

(5) Member-state implementation and governance orientation. Concrete curriculum design and instructional implementation remain matters of member-state sovereignty, and national approaches vary considerably [TBD: list of representative member-state K-12 AI curriculum policies/document numbers]. The overall governance orientation emphasizes data protection (layered atop the General Data Protection Regulation, GDPR), algorithmic transparency, minors' rights, and compliance review plus human oversight for high-risk AI used in education [11][12].

2.3.4 A supranational reference point: UNESCO's global frameworks

As a non-binding but widely influential global public good, UNESCO supplies all three jurisdictions with a shared normative vocabulary [19][21][22]:

- The **Guidance for Generative AI in Education and Research** (2023) — the world's first such guidance — sets out seven key steps for governments to regulate generative AI, advocates adopting data-protection and privacy standards and strengthening teacher training, and recommends setting **13 as the minimum age** for classroom use of generative AI tools [19][20].
- The **AI Competency Framework for Students** (released September 3, 2024, at Digital Learning Week) organizes **12 competencies** across **4 dimensions** — human-centered mindset, AI ethics, AI techniques and applications, and AI system design — progressing through three tiers: Understand, Apply, Create [21].
- The **AI Competency Framework for Teachers** (released the same period in 2024) organizes **15 competencies** across **5 dimensions** — human-centered mindset, AI ethics, AI foundations and applications, AI pedagogy, and AI for professional development — progressing through three tiers: Acquire, Deepen, Create [22].

These two frameworks supply international coordinates for this Blue Book's Chapter 3 literacy tiers and its teacher-focused chapter, and are also a key external source behind the convergence in how China, the U.S., and the EU each frame their competency goals.

2.4 Comparative Summary Table

The table below structures the comparison across the three jurisdictions (plus the UNESCO reference point) along the four dimensions. Every specific fact in the table is backed by this chapter's reference sources; anything not yet verified is marked with a placeholder.

Comparative dimension	China	United States	European Union
Dominant governance tier	Unified national deployment (top-down) [1][3]	Federal guidance + state/district autonomy (fragmented) [5][7]	Horizontal legislation + member-state implementation (layered) [11][14]
Core policy/regulatory instrument	Document No. 32 [2024], *Notice on Strengthening AI Education in Primary and Secondary Schools*; Document No. 1 [2026], *"AI+ Education" Action Plan* [1][3]	OET *Empowering Education Leaders* toolkit (Oct. 2024); White House *Advancing AI Education for American Youth* executive order (Apr. 2025) [5][6][9]	Regulation (EU) 2024/1689 (AI Act); *Digital Education Action Plan 2021–2027* [11][14]
K-12 AI curriculum standard	Progressive by grade band, standalone/embedded in parallel; essentially universal by 2030; Beijing ≥8 class hours/year from fall 2025 [1][4]	No federal standard; roughly 33–35 states have issued guidance, varying widely in force and structure [7][8]	Curriculum is member-state sovereignty; Commission supplies DigComp 2.2 and the educators' ethics guidelines

			as reference points [16][17]
Structure of literacy goals	Grade-band progression from perceive/understand/apply/create (knowledge–skills–ethics) [1]	Largely reliant on NGO frameworks such as TeachAI; no unified structure [7]	DigComp 2.2: 5 areas, 21 competencies, 250+ AI-related examples [16]
Teacher AI literacy requirements	"Teacher pipeline expansion at scale"; planned inclusion in teacher qualification exams and certification; Beijing's "Hundred-Thousand Seed Program" [3][10]	Toolkit advocates educator AI literacy (Module 8); no mandatory requirement [5]	Educators' AI and data ethics guidelines (Oct. 2022); statutory AI literacy obligation (from Feb. 2025) [11][17]
Rules on generative-AI classroom use	Mostly guidance-style language [TBD: document number of any dedicated rule]	District-level policies proliferating quickly; no national standard [7]	Emotion recognition banned in schools (Art. 5, from Feb. 2025); admissions/assessment uses are high-risk (Annex III) [11][12]
Student data and minors' protection	Minors' protection and student-data requirements [TBD: applicable regulatory provisions]	FERPA, COPPA (under 13) + roughly 40 state privacy laws [23]	GDPR + AI Act high-risk compliance [11][23]
Academic integrity and assessment governance	Regular instruction and assessment [TBD: specific provisions]	Fragmented district policies [7]	Exam-cheating monitoring AI classified as high-risk and regulated [12]
2024–2026 adoption/implementation progress	184 national-level education bases (Feb. 2024); Beijing full grade-band rollout (fall 2025) [3][4]	Roughly 33–35 states have issued guidance (2025) [7][8]	Prohibitions took effect Feb. 2025; high-risk obligations pushed to Dec. 2027 [11][15]
Global reference (UNESCO)	Student framework: 4 dimensions/12 competencies; teacher framework: 5 dimensions/15 competencies (Sept. 2024); recommended minimum age of 13 for generative AI (2023) [19][21][22]		

2.5 Convergence, Divergence, and Institutional Gaps

2.5.1 Where the three jurisdictions converge

· **The mainstreaming of AI literacy**: all three have moved AI literacy from a niche elective for a select few toward a foundational competency expected of every student, and all consistently foreground ethics and responsible use rather than purely technical training. China uses "essentially universal by 2030" as its benchmark [1]; the U.S. is driving K-12 literacy-resource development through executive action [6]; the EU has elevated AI literacy to a statutory obligation [11]; and UNESCO supplies an internationally shared set of competency benchmarks [21][22].

· **Governance and instruction treated as equally important**: as generative AI has spread, "how to use it safely" has become just as central a question as "what to teach," with data protection, minors' protection, and academic integrity emerging as shared concerns across all three. The EU folds education directly into its AI Act's prohibited/high-risk lists [11][12]; the U.S. layers state law atop FERPA/COPPA [23]; China governs through the overarching principles of moral education and minors' protection.

· **Teachers as the pivotal variable**: all three increasingly emphasize teacher AI literacy and "humans in the loop" — China plans to fold AI into teacher qualification exams and certification [3], the U.S. toolkit dedicates a module to educator literacy [5], the EU has issued educators' ethics guidelines alongside its statutory AI literacy obligation [11][17], and UNESCO has published a dedicated teacher competency framework [22]. All three treat teachers as the decisive factor in whether implementation succeeds.

2.5.2 Key divergences

· **Unified versus fragmented**: China favors a unified national timetable and deployment (the 2030 target, a system spanning all grade levels) [1][3]; the U.S. relies heavily on state/district autonomy (roughly 33–35 states each charting their own course) [7]; the EU sits in between, with horizontal legislation setting compliance boundaries while member states retain control over curriculum [11][14].

· **Strength of legislative constraint**: the EU is defined by binding horizontal regulation (prohibitions, high-risk obligations, a codified AI literacy requirement) [11]; China relies mainly on guidance-based deployment — strong in mobilizing power but not legislatively binding [1]; U.S. federal action leans toward non-mandatory guidance and incentives [6].

· **Form of curriculum delivery**: China runs "standalone courses alongside subject integration" in parallel [1]; the U.S. has no uniform form, filled in instead by states, districts, and NGO frameworks [7]; EU curriculum design is member-state sovereignty, with the Commission supplying only literacy reference points [16]. The three jurisdictions have each made a different choice.

2.5.3 Shared institutional gaps

· **Policy lagging behind agents and multimodal edge devices**: most existing texts target text-based generative AI. Although the EU's AI Act already covers emotion recognition and admissions/assessment systems [11][12], all three jurisdictions generally lack dedicated rules for autonomous agent behavior, on-device multimodal data collection, and wearable/robotic devices entering schools — a gap especially pronounced in China and the U.S. Governance needs created by these new interaction modalities remain largely unaddressed (for device-specific detail, see AI-SLI's *AI Glasses in Education White Paper 2026* and *Global Development of Educational Robots White Paper 2026*).

· **Absence of AI literacy assessment standards**: all three jurisdictions generally lack a unified, comparable K-12 AI literacy assessment instrument or attainment standard. UNESCO's frameworks supply competency dimensions and progression tiers [21][22], but remain a "competency map" rather than an "assessment rubric," constraining evidence-based evaluation (see this Blue Book's assessment and governance chapter).

· **Thin provisions on equity of access**: safeguards addressing urban–rural disparities, special-needs populations, and under-resourced regions remain largely principle-level statements in most policies, lacking operational mechanisms or investment benchmarks. U.S. compliance-gap data — 42% of districts without a signed data-processing agreement, 31% of administrators unable to identify the applicable law — indirectly confirms the gulf between implementation capacity and policy text [23] (see this Blue Book's dedicated equity-of-access chapter).

2.6 Summary

Between 2024 and 2026, China, the United States, and the EU each completed a shared shift in their K-12 AI education policies — from supplying "informatization materials" to balancing capability-building with governance — converging strongly on three points: the mainstreaming of AI literacy, the pivotal role of teachers, and equal weight given to governance and instruction. Each jurisdiction's flagship instruments — China's Document No. 32 [2024] and Document No. 1 [2026] [1][3]; the U.S. OET toolkit and the 2025 White House executive order [5][6]; the EU's Regulation (EU) 2024/1689 and its Digital Education Action Plan [11][14] — respond to the same governance challenge through three distinct paradigms: unified deployment with a quantified timetable; federal incentives paired with state/district autonomy; and horizontal legislation paired with a codified literacy obligation. The main divergences lie in how unified versus fragmented governance is, how strong the legislative constraint is, and whether curriculum delivery is standalone or embedded. The institutional gaps shared by all three center on lagging governance of agents and multimodal edge devices, the absence of AI literacy assessment standards, and thin operational provisions for equity of access. These three gaps define the problem space that later chapters of this Blue Book — on literacy frameworks, human–AI collaboration, assessment, governance, equity of access, and teacher literacy — take up in turn. Any specific fact in this chapter not yet confirmed by a verified source (such as China's dedicated rules on generative AI entering schools, or the list of representative EU member-state curriculum policies) remains marked with a `[TBD: ...]` placeholder and must not be filled in with unverified figures or proper nouns.

Chapter References

[1] Notice on Strengthening Artificial Intelligence Education in Primary and Secondary Schools (Document No. 32 [2024]) [关于加强中小学人工智能教育的通知（教基厅函〔2024〕32 号）]. General Office of the Ministry of Education of the People's Republic of China. 2024. http://www.moe.gov.cn/jyb_xwfb/gzdt_gzdt/s5987/202412/t20241202_1165500.html

[2] Ministry of Education Deploys Measures to Strengthen AI Education in Primary and Secondary Schools (policy interpretation, covering overall requirements, the six major tasks, and grade-band divisions) [教育部部署加强中小学人工智能教育（部署解读，含总体要求与六大任务、学段划分）]. Ministry of Education of the People's Republic of China. 2024. http://www.moe.gov.cn/jyb_xwfb/gzdt_gzdt/s5987/202412/t20241202_1165500.html

[3] Notice of the Ministry of Education and Four Other Departments on Issuing the "AI+ Education" Action Plan (Document No. 1 [2026]) [教育部等五部门关于印发《"人工智能+教育"行动计划》的通知（教科信〔2026〕1 号）]. Ministry of Education of the People's Republic of China and four other departments. 2026. http://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html

[4] Notice of the Beijing Municipal Education Commission on Issuing the Beijing Local Curriculum Outline for AI Education in Primary and Secondary Schools (Trial) (2025 Edition) [北京市教育委员会关于印发《北京市中小学人工智能教育地方课程纲要（试行）（2025 年版）》的通知]. Beijing Municipal Education Commission. 2025. https://jw.beijing.gov.cn/xxgk/2024zcwj/2024qtwj/202506/t20250626_4127392.html

[5] Empowering Education Leaders: A Toolkit for Safe, Ethical, and Equitable AI Integration. U.S. Department of Education, Office of Educational Technology. 2024. https://eric.ed.gov/?id=ED661924

[6] U.S. Department of Education Issues Guidance on Artificial Intelligence Use in Schools. U.S. Department of Education. 2024. https://www.ed.gov/about/news/press-release/us-department-of-education-issues-guidance-artificial-intelligence-use-schools-proposes-additional-supplemental-priority

[7] State AI Guidance for Education. AI for Education. 2025. https://www.aiforeducation.io/ai-resources/state-ai-guidance

[8] 31 States Have Adopted AI Education Policy Guidelines. AI for Education. 2025. https://www.aiforeducation.io/blog/31-states-adopt-state-ai-guidlines

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

[10] Primary and Secondary Schools to Offer General AI Course Starting Fall Semester, No Fewer Than 8 Class Hours Per Year, Spanning Primary Through High School [中小学秋季学期起开设 AI 通识课 每学年不少于 8 课时 覆盖小学至高中全学段]. Beijing Municipal People's Government Portal (Capital Window) [首都之窗]. 2025. https://www.beijing.gov.cn/ywdt/gzdt/202503/t20250308_4029252.html

[11] Regulation (EU) 2024/1689 (Artificial Intelligence Act). European Parliament and Council. 2024. https://artificialintelligenceact.eu/high-level-summary/

[12] Annex III: High-Risk AI Systems (including education and vocational training scenarios). EU Artificial Intelligence Act. 2024. https://artificialintelligenceact.eu/annex/3/

[14] Digital Education Action Plan (2021–2027). European Commission / EUR-Lex. 2020. https://eur-lex.europa.eu/EN/legal-content/summary/digital-education-action-plan-2021-2027.html

[15] Implementation Timeline (AI Act phased application and the Digital Omnibus delay). EU Artificial Intelligence Act. 2024–2026. https://artificialintelligenceact.eu/implementation-timeline/

[16] DigComp 2.2: The Digital Competence Framework for Citizens — With new examples of knowledge, skills and attitudes. European Commission, Joint Research Centre (JRC). 2022. https://publications.jrc.ec.europa.eu/repository/handle/JRC128415

[17] Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators. European Commission. 2022. https://education.ec.europa.eu/focus-topics/digital-education/actions/plan/ethical-guidelines-for-educators-on-using-artificial-intelligence

[18] Ethical guidelines on the use of AI and data in teaching and learning for educators (full PDF). European Commission. 2022. https://teachertaskforce.org/sites/default/files/2023-03/2022_EC_ethical%20guidelines%20on%20the%20use%20of%20artificial%20intelligence_EN.pdf

[19] Guidance for generative AI in education and research. UNESCO. 2023. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

[20] UNESCO: Governments must quickly regulate Generative AI in schools (including the recommended minimum age of 13 and the seven key steps). UNESCO. 2023. https://www.unesco.org/en/articles/unesco-governments-must-quickly-regulate-generative-ai-schools

[21] AI competency framework for students (4 dimensions / 12 competencies / 3 tiers). UNESCO. 2024. https://www.unesco.org/en/articles/ai-competency-framework-students

[22] AI competency framework for teachers (5 dimensions / 15 competencies / 3 tiers). UNESCO. 2024. https://www.unesco.org/en/articles/ai-competency-framework-teachers

[23] Overview of Federal Regulations Related to Artificial Intelligence (covering FERPA, COPPA, and the state-level privacy law landscape). National Education Association (NEA). 2025. https://www.nea.org/sites/default/files/2025-06/5.1-ai-policy-overview-of-federal-regulations-final.pdf

Chapter 3 AI Literacy Frameworks: What Students and Teachers Need to Know

This chapter builds on the AI Competency Framework for Students and the AI Competency Framework for Teachers, both released by UNESCO in August 2024 [1][2]. Against the 2026 reality in which generative AI (GenAI) and agents have become part of everyday K-12 teaching, we reconstruct what "AI literacy" means, how it is structured, and how it develops — and we offer an analytical framework that national curricula and teacher professional standards can plug into. Working from a competency-and-governance angle, the chapter answers four questions: what does K-12 AI literacy actually consist of; how many levels does it have; how should it be built jointly on the student and teacher sides; and how are major economies writing it into policy?

Fig. 4 Structure of UNESCO's 2024 student/teacher AI competency frameworks

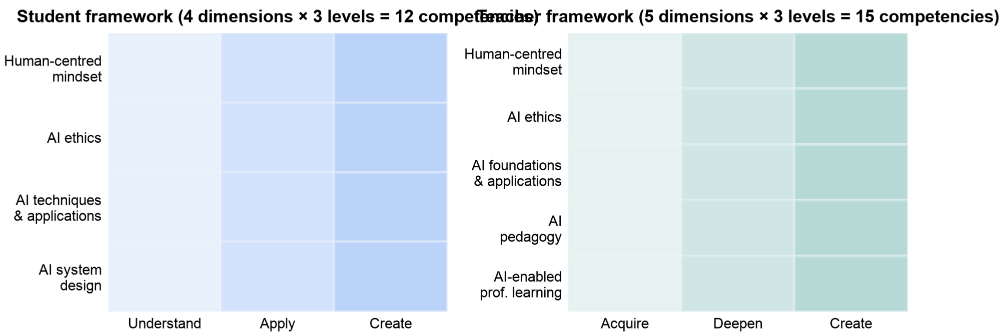

Source: UNESCO, AI Competency Frameworks for Students and Teachers (2024).

3.1 From "knowing how to use the tool" to a competency-values compound: redefining AI literacy

Before generative AI and agents became mainstream, both scholarship and policy tended to narrow "AI literacy" down to an extension of digital skills — recognizing AI, operating a handful of applications. The 2022 parent white paper reflected exactly this framing, organizing its content around material categories such as resources, teaching materials, algorithms, and labs, with the implicit assumption that literacy meant simply having been exposed to the relevant resources. That assumption was demolished by hard data in 2023. A global survey UNESCO conducted in May 2023 for the Ministerial Roundtable on generative AI and education covered more than 450 schools and universities. It found that **fewer than 10% of educational institutions had adopted an institutional policy or formal guidance on generative AI use** — roughly 13% among universities and only about 7% among schools — and that of the institutions claiming to have guidance, about 40% had communicated it only verbally, with nothing in writing [3]. In other words, in the first year that

ChatGPT reached classrooms, the world's education systems were, institutionally speaking, caught almost entirely unprepared. It was precisely this governance vacuum that triggered the dense wave of frameworks issued between 2023 and 2026.

Two structural shifts since 2023 have made the old "literacy equals knowing how to use the tool" definition insufficient:

- **The bar for capability has dropped, and its boundary has expanded.** GenAI turned tasks that once required specialized training — generating text, images, and code — into everyday actions triggered by natural language alone. Literacy's center of gravity has shifted from "can you use it" to "when should you use it, how do you judge it, can you trust it, and who is accountable."
- **The human–machine relationship has moved from tool use to collaboration and delegation.** Agents can now carry out multi-step tasks autonomously on a person's behalf. Students and teachers are no longer dealing with a passive tool but with an "agentic system" that requires oversight, verification, and constraint. Literacy must therefore include a critical understanding of AI's limitations, biases, hallucinations, and the boundaries of responsibility.

Building on this, this Blue Book adopts the following working definition: **K-12 AI literacy is a compound competency unifying knowledge, skills, values, and action — the integrated capacity of learners and educators to understand, use, evaluate, and jointly govern AI systems safely, critically, responsibly, and creatively in real-world contexts.** This definition deliberately elevates values and ethics from a subordinate item to a pillar standing alongside knowledge and skills, echoing the overarching orientation shared by both UNESCO frameworks — each places a "human-centred mindset" as its first dimension and positions students as "AI co-creators and responsible citizens" rather than passive users [1][4]. This orientation is structurally near-identical to the "four-in-one AI literacy" put forward in China's Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition), which fuses knowledge, skills, thinking, and values [5].

> For how the underlying trends toward on-device processing, multimodality, and agentic behavior are entering the classroom, see the scenario analyses in this institute's Global AI Glasses in Education Blue Book 2026 and Global Educational Robotics Development White Paper 2026.

3.2 The Student AI Competency Framework: dimensions and progression

UNESCO's 2024 AI Competency Framework for Students was authored by Miao Fengchun, Kelly Shiohira, and Natalie Lao, and released on August 8, 2024, alongside Digital Learning Week 2024 [1][4]. The framework is organized on a two-axis structure — competency dimension by progression level — **setting out 12 competencies across four dimensions, with each competency spanning three progression levels** [1][4]. This section follows the framework's published dimensional architecture. The verbatim wording of each competency item under every dimension, along with the recommended grade bands and sample task lists, should be taken from the framework's original text;

wherever this Blue Book has not been able to verify wording against the primary source directly, it is marked **[TBD]**, and no item name has been invented [TBD: the verbatim official wording of the 12 competencies in UNESCO's 2024 student framework, and their grade-band mapping, per the original text].

3.2.1 Four competency dimensions

The framework groups student AI competencies into four mutually reinforcing dimensions, each comprising 3 competencies, for a total of 12 [1][4]:

- **A human-centred mindset**: Places human agency, rights, and well-being above AI, understanding AI as a means serving people rather than an end in itself; emphasizes that students should establish and assert their own agency in their relationship with AI, and maintain a critical, questioning stance toward automated decisions [4][6].
- **Ethics of AI**: Recognizing and responding to ethical issues such as fairness, privacy, transparency, accountability, and safety; advocating "responsible and safe use" and "ethics-by-design," and exercising responsible value judgment while using AI [4][6].
- **AI techniques and applications**: Understanding the fundamental principles of AI — data, algorithms, models, training, and inference — laying a foundational base of AI knowledge and skills for lifelong learning, and being able to use the appropriate tools and methods effectively on suitable tasks [1][4].
- **AI system design**: Participating in conceiving and creating AI solutions through computational and engineering thinking, from problem definition through data and solution design to testing and iteration, cultivating problem-solving, creativity, and design thinking oriented toward inclusive, sustainable design [1][4].

Note: The above reflects the four dimensions' official names and their qualitative descriptions as given in the framework. The **relative weighting, recommended instructional hours, and share of assessment** across the four dimensions are not strongly specified by the framework; where quantification is needed for curriculum implementation, national curriculum documents should be the authority [TBD: instructional-hour or credit allocations for the four dimensions in Chinese/U.S./European curricula].

3.2.2 Progression levels

The framework sets a three-level progression, shallow to deep, for every competency, officially named **Understand — Apply — Create**, used to distinguish the depth expected of students at different grade levels and with different degrees of experience [1][4]. The specific performance descriptors at each level should be taken from the original text [TBD: the verbatim performance descriptors for each level in UNESCO's 2024 student framework]. To facilitate cross-national curriculum alignment, this Blue Book summarizes the framework's own three-level naming below and leaves a column open for mapping to each country's existing competency-level systems:

Progression level (official framework name)	**Typical student performance (qualitative description)**	**Recommended grade band**	**Mapping to national curriculum levels**
Understand	Can explain basic AI concepts, common applications, and potential risks	[TBD: grade-band recommendation]	[TBD: mapping]
Apply	Can select and use AI tools appropriately, under guidance, to complete learning tasks, and verify the results	[TBD: grade-band recommendation]	[TBD: mapping]
Create	Can design a simple AI solution, or offer a well-grounded critique of AI outputs and their social impact	[TBD: grade-band recommendation]	[TBD: mapping]

The "recommended grade band" and "curriculum-level mapping" columns in the table above must be filled in against each country's actual policy; this Blue Book does not invent correspondences. It is worth noting that UNESCO's three-level Understand–Apply–Create progression can be conceptually cross-referenced with China's three-stage spiral — primary school (experience and interest), middle school (problem-solving), high school (innovative application) — as well as with the EU–OECD framework's competency spectrum of Engage–Create–Manage–Shape (see the comparison in Section 3.6) [1][5][7].

3.3 New core sub-competencies for the generative-AI era

The two UNESCO frameworks provide the overarching dimensions, but putting GenAI and agents into practice at the K-12 level demands further specification of several **new sub-competencies**. This section treats these as contextual elaborations of the frameworks rather than replacements for them; every sub-competency below can be traced back to one of the existing framework dimensions, forming a mapping from "framework dimension" to "contextualized sub-competency."

3.3.1 Prompting & interaction literacy

Understanding that GenAI generates content probabilistically, and mastering the interactive methods of stating tasks clearly, supplying constraints and context, and iterating through step-by-step follow-up questions and revisions — while also understanding that "prompt engineering" is not a cure-all, since model capability and training data carry inherent limits. This sub-competency corresponds to the framework's "AI techniques and applications" dimension.

3.3.2 Critical verification

Treating "AI can confidently give wrong answers (hallucination)" as a default assumption, and building the habit of independently verifying sources, facts, citations, and logic; understanding that generated content is not equivalent to trustworthy evidence. UNESCO's student framework explicitly lists "critical judgement of AI solutions" as a core orientation [4]; China's Guidelines on the Use of Generative AI by Primary and Secondary School Students (2025 Edition) likewise requires "guiding students to cross-verify the plausibility of generated content" and explicitly prohibits "directly copying AI-generated content as homework or exam answers" [8][9]. The related evidence-based verification methods are developed further in the chapter on assessment and academic integrity (see Chapter [TBD] of this Blue Book).

3.3.3 Agentic collaboration

Understanding that agents can autonomously execute tasks across multiple steps, and mastering a "delegate–supervise–take over" collaboration pattern that makes clear which tasks may be delegated and which decisions and responsibilities must remain with a human. China's Guidelines already state explicitly, at the high-school stage, that students should "develop cross-disciplinary integration solutions based on agentic tools," formally incorporating agentic collaboration into grade-band objectives [5]. This sub-competency also touches on the human-accountability requirements embedded in both the "human-centred mindset" and "ethics of AI" dimensions.

3.3.4 Data & privacy

Understanding how personal data flows and what risks arise during AI interactions, and mastering the basic norms of minimal disclosure, anonymization, and compliant use. China's Guidelines set an operational red line at the institutional level: **students and teachers are strictly prohibited from entering exam questions, personal identifying information, or other sensitive data when using generative AI tools, and schools are required to establish a "whitelist" system for generative AI tools** [8][9]. This sub-competency corresponds to the framework's "ethics of AI" dimension. [TBD: the specific document numbers and scope of application of the cited regulations on minors' data protection, such as each country's personal information protection law].

3.3.5 Ethics & fairness

Recognizing bias in training data, algorithmic discrimination, and gaps in access, and understanding the differentiated impact AI may have on different groups (urban/rural, language, disability, and others). The UNESCO frameworks weave "ethics-by-design" and "inclusive, sustainable AI design" throughout [4]; the EU–OECD framework, for its part, sets "Shape AI" apart as a distinct competency domain, pointing to students' active engagement with AI's social and governance impact

[7]. This topic is mutually reinforcing with the dedicated chapter on equity of access (see Chapter [TBD] of this Blue Book).

3.4 The Teacher AI Competency Framework: from user to designer and gatekeeper

UNESCO's 2024 AI Competency Framework for Teachers was authored by Miao Fengchun and Mutlu Cukurova, and released the same day as the student framework — August 8, 2024 [2][4]. It shares its origin with the student framework but differs in emphasis: teachers are not merely users of AI but the **designers, models, and ethical gatekeepers** of AI use in the classroom. The framework follows the same "competency dimension × progression level" structure, **setting out 15 competencies across five dimensions, with each competency spanning three progression levels** [2][4].

3.4.1 Teacher competency dimensions

The framework organizes its five dimensions around the professional context of teaching; their official names and qualitative emphases are as follows [2][4][10]:

- **Human-centred mindset**: Upholding a human-centred stance, adopting a reflective and ethical attitude toward AI, and modeling respect for student rights and responsible AI use.
- **Ethics of AI**: Addressing privacy, data protection, and the avoidance of algorithmic bias, and putting compliance and fairness into practice in both teaching and administration.
- **AI foundations and applications**: Building a technical understanding of how AI works, and being able to evaluate, select, and safely apply AI within educational contexts.
- **AI pedagogy**: Using AI to enhance lesson preparation, instruction, and assessment, and embedding AI effectively into instructional design and differentiated support rather than simply using it to replace direct instruction.
- **AI for professional learning**: Using AI to support teachers' own continuous learning and professional collaboration, and participating in the sharing of experience.

Note: The official names of the five dimensions are given above. The verbatim wording of the 15 competency items under each dimension, and their formal correspondence with existing teacher digital/ICT competency frameworks (such as UNESCO's ICT Competency Framework for Teachers), should be taken from the framework's original text; this Blue Book does not invent them [TBD: the verbatim official wording of the 15 competencies in UNESCO's 2024 teacher framework; the original-text correspondence with existing teacher digital/ICT competency frameworks]. To facilitate alignment with teacher professional standards, this Blue Book regroups the five dimensions above into four practice-oriented lines organized around the teacher's actual workflow — values and ethical stance, AI foundations and tool command, pedagogical integration, and professional development and shared governance. This regrouping is an analytical restatement and does not alter the framework's original five-dimension official structure.

3.4.2 Teacher progression levels

The framework sets a three-level progression for teacher competency, officially named **Acquire — Deepen — Create**, charting a development path from initial exposure to the ability to lead peers and innovate teaching through AI [2][4][10]. The behavioral descriptors for each level should be taken from the original text [TBD: the verbatim performance descriptors for each level in UNESCO's 2024 teacher framework].

Progression level (official framework name)	Typical teacher performance (qualitative description)	Supporting professional-development measures
Acquire	Has foundational AI literacy, understands how AI works, and attempts to use it with support	[TBD: training/professional-development measures]
Deepen	Can integrate AI into teaching practice, critically evaluate its effects, and handle common ethical and compliance issues	[TBD: training/professional-development measures]
Create	Can innovate teaching through AI, design new methods, mentor peers, and participate in shaping governance rules	[TBD: training/professional-development measures]

The proportion of teachers at each level, the standards for reaching each level, and the required professional-development hours must be filled in against real survey data and policy documents [TBD: survey data on the current state of teacher AI literacy and professional-development requirements].

3.4.3 The real-world gap in teacher AI literacy

Building teacher capacity is the binding constraint on putting AI literacy into practice, and this constraint is clearly visible in the data. According to U.S. data cited by the TeachAI coalition, as of April 2025 only 26 states nationwide had issued official guidance on K-12 AI use; only 18% of principals reported that their school or district provided guidance on AI use, and that share fell to just 13% in high-poverty schools, well below the 25% figure at wealthier schools [11]. UNESCO's 2023 global survey corroborates this: formal institutional-level governance stood at under 10%, and much of what existed was verbal rather than written [3]. The gap concentrates in three areas: **lagging pre-service preparation** (teacher-training programs have yet to systematically incorporate AI pedagogy), **uneven in-service professional-development coverage**, and **divergence in capacity and institutional arrangements between urban and rural areas and across schools**. The specific figures, teacher self-assessments, and needs-survey findings are developed in the dedicated chapter on teacher literacy (see Chapter [TBD] of this Blue Book); this chapter does not pre-populate unverified figures [TBD: local survey data and sources on teacher AI literacy coverage rates, training participation rates, and urban–rural disparities].

3.5 Coupling student and teacher competency: a double-helix model

Student literacy and teacher literacy are not two parallel checklists — they form a coupled system that mutually constructs and mutually constrains. The very fact that UNESCO released two "sister frameworks" on the same day, using the same methodology, reflects the underlying logic that the two must be built in coordination [4]. Building on this, this Blue Book proposes a "double-helix" analytical framework to characterize this coordinating mechanism (this model is an analytical construct for organizing the narrative and carries no empirical parameters):

- **Modeling transmission (teacher → student)**: A teacher's value stance and norms of use are internalized by students, through daily classroom experience, into tacit habits of AI use. The teacher framework places "human-centred mindset" as its first dimension precisely so that teachers become models of responsible use in the first place [2].
- **Demand pull (student → teacher)**: Student literacy goals — such as critical verification and agentic collaboration — in turn define the "AI pedagogy" competency teachers must hold. This is exactly the dimension the student framework lacks and the teacher framework alone provides [1][2].
- **Shared governance (teachers and students → school-level → regional/national)**: Rules on academic integrity, data privacy, and the like are jointly practiced by teachers and students in real settings, forming school-level governance conventions — such as the tool "whitelist" system required by China's Guidelines [8] — that then connect upward to regional and national policy (see the governance chapter of this Blue Book, Chapter [TBD]).

The model yields a strong proposition: **any approach that builds literacy requirements only on the student side, without simultaneously building teacher capacity and school-based governance, will leave literacy goals unanchored.** The data gap described in Section 3.4.3 — insufficient guidance coverage, urban–rural divergence — is exactly the empirical symptom of a weak "teacher arm" in the double helix.

3.6 A framework for curriculum implementation: four frameworks compared side by side

Between 2024 and 2026, the world's major K-12 AI literacy frameworks moved from scarcity to a multi-source landscape. This section places the two UNESCO frameworks alongside the two other authoritative frameworks, extracting comparable dimensions to provide a common base for the chapters that follow on Chinese, U.S., and European curricula. The key facts of the four frameworks are compared below:

Framework	Publisher	Date	Target audience	Structure (dimensions × levels / item count)	Progression/competency-domain naming
AI Competency Framework for Students	UNESCO	2024-08	K-12 students	4 dimensions × 3 levels, 12 competencies [1][4]	Understand—Apply—Create
AI Competency Framework for Teachers	UNESCO	2024-08	K-12 teachers	5 dimensions × 3 levels, 15 competencies [2][4]	Acquire—Deepen—Create
Empowering Learners for the Age of AI (AI Literacy Framework)	European Commission + OECD (with CodeAI support)	Review draft 2025-05 / formally released 2026-06	K-12 students	4 competency domains, 19 competencies, organized by knowledge–skills–attitudes [7][12]	Engage—Create—Manage—Shape
Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)	Teaching Guidance Committee for Basic Education, China's Ministry of Education	2025-05-13	K-12 students	Three-grade-band spiral, four-in-one literacy [5]	Primary—Middle—High school

A few evidence-based observations:

1. **"Human-centeredness" has become the largest global common denominator.** Both UNESCO frameworks place "human-centred mindset" as their first dimension [1][2]; the EU–OECD framework treats "knowledge–skills–attitudes" as the shared foundation running through every competency and sets "Shape AI" apart as its own domain, emphasizing students' active engagement with AI's social impact [7]; China's Guidelines include "AI literacy and social-responsibility awareness" within its four-in-one structure [5]. The wording differs across the three, but the underlying values converge.
2. **Assessment is catching up.** The EU–OECD framework explicitly connects to the **PISA 2029 "Media and AI Literacy" assessment** [12], meaning AI literacy will, for the first time, enter the

scales used in large-scale international learning assessments — which will in turn push national curricula from advocacy toward measurability.

3. **The distinctiveness of China's approach lies in pairing norms with literacy.** Alongside the general-education guidelines, China simultaneously issued the Guidelines on the Use of Generative AI by Primary and Secondary School Students (2025 Edition), which sets usage boundaries by grade band: **primary-school students must use open-ended content-generation features only appropriately and with help from teachers or parents, and may not use them unsupervised; middle-school students may explore, to an appropriate degree, the logical analysis of generated content; high-school students may undertake inquiry-based learning that draws on the underlying technical principles** — with clear red lines on data handling and academic integrity [8][9]. This two-document structure — "competency framework plus usage norms" — is the most complete of the four frameworks.

Building on this comparison, subsequent curriculum chapters can populate the template below cell by cell with each country's actual policy text, ensuring cross-national comparability (cell data must be filled in from real policy texts; this Blue Book does not invent any document number, year, or course name):

Alignment dimension	China	United States	European Union
Policy positioning of AI literacy and its documentary basis	Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition); Guidelines on the Use of Generative AI by Primary and Secondary School Students (2025 Edition), Teaching Guidance Committee for Basic Education, Ministry of Education, 2025-05-13 [5][8]	[TBD: federal/state-level document names and issuing bodies, such as individual state AI guidance]	Empowering Learners for the Age of AI (European Commission + OECD); Ethical guidelines on the use of AI and data in teaching and learning for educators (2022) [7][13]
Standalone course versus integration into existing subjects	Standalone courses, cross-disciplinary integration, and practical activities "flexibly adopted" [5]	[TBD]	[TBD: mostly cross-disciplinary integration]
Grade bands covered and starting grade	Three-grade-band spiral across primary, middle, and high school [5]	[TBD]	Primary and secondary [7]
Correspondence with this Blue Book's four dimensions	Four-in-one (knowledge/skills/thinking/values) [5]	[TBD]	4 domains: Engage/Create/Manage/Shape [7]
Teacher literacy requirements and professional-	[TBD: national teacher professional-development	As of 2025-04, only 26 states had issued AI	Ethical guidelines on AI use for educators (2022, including

development arrangements	requirements]	guidance [11]	a GenAI revision) [13]
Assessment and academic-integrity provisions	Prohibits directly copying AI-generated content for homework/exams; tool whitelist; data red lines [8][9]	[TBD]	Connects to the PISA 2029 Media and AI Literacy assessment [12]

3.7 Chapter summary

Against the 2026 context of generative AI and agentic behavior, this chapter reconstructed the definition of K-12 AI literacy, establishing it as a human-centered compound competency unifying knowledge, skills, values, and action. Drawing on the "dimension × progression" architecture of UNESCO's 2024 student framework (4 dimensions, 12 competencies, Understand–Apply–Create) and teacher framework (5 dimensions, 15 competencies, Acquire–Deepen–Create) [1][2], the chapter mapped the competency dimensions and progression pathways on both sides, and supplemented them with new GenAI-era sub-competencies — prompting and interaction, critical verification, agentic collaboration, data and privacy, and ethics and fairness — each traced back to its originating dimension in the parent frameworks. The chapter then proposed a "double-helix" coupling model showing that student literacy and teacher literacy must be built in coordination, and must be mutually reinforced by school-based governance; the current-state data from TeachAI and UNESCO — guidance coverage under 10%, only 26 U.S. states with guidance issued, and urban–rural/inter-school divergence — is exactly the empirical evidence of a weak "teacher arm" [3][11]. Finally, the chapter placed the two UNESCO frameworks side by side with the EU–OECD's Empowering Learners for the Age of AI (4 competency domains, 19 competencies, aligned with PISA 2029) and China's Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [5][7][12], distilling three judgments: human-centeredness has become a global common denominator, assessment is accelerating to catch up, and China's approach pairs norms with literacy. The chapter also provided a template for cross-national curriculum alignment. Wherever specific figures, document numbers, verbatim item names, or survey findings could not be verified against primary sources, this chapter has marked them **[TBD: …]**, leaving them for evidence-based completion rather than inventing any of them.

Chapter References

[1] AI competency framework for students (学生人工智能能力框架) · UNESCO · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-students

[2] AI competency framework for teachers (教师人工智能能力框架) · UNESCO · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-teachers

[3] UNESCO survey: Less than 10% of schools and universities have formal guidance on AI · UNESCO · 2023 · https://www.unesco.org/en/articles/unesco-survey-less-10-schools-and-universities-have-formal-guidance-ai

[4] What you need to know about UNESCO's new AI competency frameworks for students and teachers · UNESCO · 2024 · https://www.unesco.org/en/articles/what-you-need-know-about-unescos-new-ai-competency-frameworks-students-and-teachers

[5] Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [《中小学人工智能通识教育指南（2025 年版）》] · Teaching Guidance Committee for Basic Education, Ministry of Education · 2025 · https://www.cernet.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[6] The UNESCO Competency Framework for Students · Oide Technology in Education · 2024 · https://www.oidetechnologyineducation.ie/ai-hub/projects-research-and-insights/ai-frameworks-and-guidance-for-education/the-unesco-competency-framework-for-students/

[7] AI Literacy Framework for Primary & Secondary Education (Empowering Learners for the Age of AI) · European Commission / OECD / CodeAI · 2025–2026 · https://ailiteracyframework.org/

[8] Release of the Guidelines on the Use of Generative AI by Primary and Secondary School Students (2025 Edition) [《中小学生成式人工智能使用指南（2025 年版）》发布] · China Education Online (中国教育在线) · 2025 · https://www.eol.cn/zhengce/wenjian/202505/t20250512_2667831.shtml

[9] Ministry of Education releases the Guidelines on the Use of Generative AI by Primary and Secondary School Students (2025 Edition): homework may not copy AI-generated content [教育部发布《中小学生成式人工智能使用指南（2025 年版）》作业禁止复制人工智能生成内容] · Guangming Online (光明网) · 2025 · https://tech.gmw.cn/2025-05/14/content_38025442.htm

[10] Summary of UNESCO AI Competency Framework for Teachers · CIDDL · 2024 · https://ciddl.org/summary-of-unesco-ai-competency-framework-for-teachers/

[11] TeachAI | AI Guidance for Schools Toolkit (including U.S. state-level AI guidance and principal survey data) · TeachAI (Code.org, ETS, ISTE, and others) · 2024–2025 · https://www.teachai.org/toolkit

[12] EU and OECD launch AI literacy framework for schools (including the 19 competencies and alignment with PISA 2029) · Digital Watch Observatory · 2026 · https://dig.watch/updates/eu-oecd-ai-literacy-framework-schools

[13] Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators · European Commission (Digital Education Action Plan) · 2022 · https://education.ec.europa.eu/focus-topics/digital-education/actions/plan/ethical-guidelines-for-educators-on-using-artificial-intelligence

Chapter 4 AI-Native Curriculum: Rebuilding Course Sequences by Grade Band and the Rise of Education-Vertical Large Language Models

Fig. 5 AI-native curriculum: a spiral progression across competency dimensions and school stages

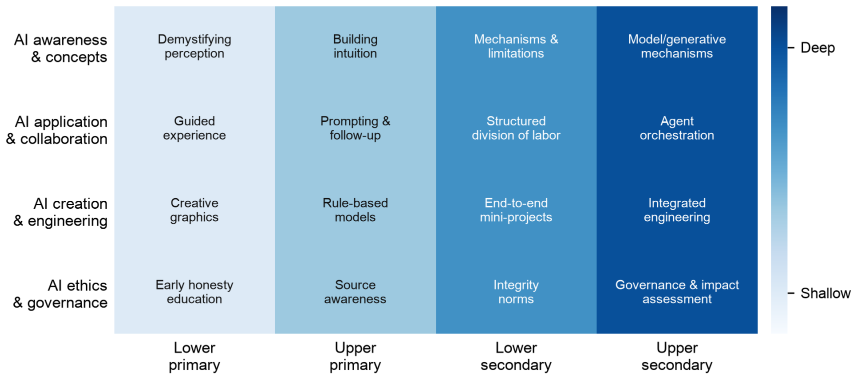

Source: this report's analytical framework (a structured re-expression of the policy and literacy frameworks in Chapters 3-4, not external empirical data).

4.1 From "Bolt-On Tool" to "AI-Native": A Generational Shift in Curriculum Design

The previous generation of K-12 AI curricula was built on a "bolt-on tool" logic: artificial intelligence was treated as a knowledge domain or a category of lab equipment, appended to existing information-technology courses and organized around a three-step arc of "learn the concept, operate the platform, complete the project." That logic made sense before generative AI and agents became widely accessible — when students had little hands-on access to capable models, it was natural for coursework to center on algorithmic principles and tool mechanics.

Since late 2022, as large language models, multimodal models, and agent technologies have moved into everyday reach, the premises underlying curriculum design have shifted structurally. Three converging pressures toward "nativeness" have emerged:

- **Capability nativeness**: The general-purpose models students already carry in their pockets can generate, retrieve, reason, and orchestrate in combination. Curricula can no longer assume "students don't know how to use AI" — they must instead assume "students are already using it, and often using it poorly."
- **Process nativeness**: AI may now participate throughout the entire arc of acquiring knowledge, creating work, and solving problems. Curricula need to make the human–AI collaborative process itself — not just the final product — an object of teaching and assessment.

- **Supply-side nativeness**: Education-vertical large language models and agent-based teaching assistants are entering the classroom supply chain. Curriculum content is now coupled to the capabilities of the underlying model, so course design must address "what to teach" and "which model to teach it with" simultaneously.

This is not merely a theoretical shift — it has been borne out by a wave of policy issued between 2024 and 2026. In November 2024, the General Office of China's Ministry of Education issued the Notice on Strengthening Artificial Intelligence Education in Primary and Secondary Schools (Document No. 32 [2024] of the General Office of Basic Education [Chinese: 教基厅函〔2024〕32号]), which called for "essentially universal AI education in primary and secondary schools before 2030" and placed "building a systematic curriculum system" first among six major tasks [1][2]. In May 2025, the Ministry's Basic Education Teaching Guidance Committee formally released the Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) alongside the companion Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition). The former called for "building a tiered, spiraling AI general-education system for primary and secondary schools," while the latter drew explicit boundaries around whether and how students may use generative AI for specific tasks [3][4]. By April 2026, the Ministry of Education and four other bodies (the National Development and Reform Commission, the Ministry of Industry and Information Technology, the Ministry of Science and Technology, and the National Data Administration) jointly issued the "AI+Education" Action Plan (Document No. 1 [2026] of the Ministry of Education's Department of Science, Technology and Informatization [Chinese: 教科信〔2026〕1号]), which went further still, calling for "developing AI education LLMs by education stage" — elevating model supply from an industry talking point to a nationally organized research priority [5][6]. Read together, these three documents — from "setting up courses," to "establishing rules for using AI," to "developing education LLMs by grade band" — trace the exact policy trajectory of the shift from tool-embedded to AI-native curricula.

Building on this trajectory, this chapter defines an **AI-native curriculum** as: **a form of K-12 curriculum organized around AI-literacy objectives, treating the human–AI collaborative process as the central object of instruction, defaulting to education-vertical models and agents as infrastructure, and systematically rebuilt as a progression across grade bands.** This definition connects directly to the literacy framework set out in Chapter 3 of this Blue Book, forming the "literacy-to-curriculum" pipeline.

> AI-SLI's view: An AI-native curriculum is not "one more AI course bolted on" — it is a re-asking of the questions existing courses already pose, once AI is part of the environment. It simultaneously changes what students learn (content), how they learn it (process), and what they learn with (infrastructure). Policy has already anchored these three dimensions separately — in the general-education guidelines, the usage

guidelines, and the push to develop education LLMs — and curriculum designers need to work all three into a single coherent picture rather than treating "the AI course" as a standalone project.

4.2 A Two-Dimensional Framework for Rebuilding the Curriculum by Grade Band

K-12 education spans multiple stages of cognitive development, so a one-size-fits-all AI curriculum fits neither how students learn nor the constraints of limited class time. This section proposes a rebuilding framework organized along two axes — **grade band and capability dimension**. The capability dimension carries over the literacy structure from Chapter 3, divided into four strands: AI awareness and concepts, AI application and collaboration, AI creation and engineering, and AI ethics and governance literacy. This four-strand approach mirrors the Guidelines for General AI Education (2025 Edition), which calls for "an integrated fusion of knowledge, skills, thinking, and values" to form a "four-in-one" AI literacy [3]. The table below sets out how the emphasis and depth of each strand progresses by grade band.

Grade Band	AI Awareness & Concepts	AI Application & Collaboration	AI Creation & Engineering	AI Ethics & Governance Literacy
Lower primary	Perceiving "intelligence" phenomena; demystifying anthropomorphism	Teacher-guided, experiential use	No-code/visual creative expression	Honest use; early privacy awareness
Upper primary	Building intuition for data–pattern–prediction	Basic prompt construction and follow-up questioning	Simple datasets and rule-based models	Information discernment; source awareness
Middle school	Basic mechanics and limitations of machine learning	Structured prompting; designing human–AI division of labor	End-to-end mini-projects; calling models	Bias, copyright, and academic-integrity norms
High school	Model principles; generative and agentic mechanisms	Orchestrating agents for complex tasks	Comprehensive, cross-disciplinary engineering applications	Governance frameworks; social-impact assessment

This progression tracks closely with the official framing in the Guidelines for General AI Education (2025 Edition), which positions primary, middle, and high school respectively as "building experience and interest" (awareness), "understanding the logic of the technology" (awareness), and "technology strategy" (awareness) — with the skills strand moving in turn through "basic application ability," "solving real-world problems," and "innovative application"; the thinking strand rising from "foundational thinking" through "engineering thinking" to "systems thinking"; and the values strand progressing from "cultural awareness and safety habits" through "deepened ethical awareness" to "practicing social responsibility" [3]. The 2024 Notice expressed the same progression in its own terms: "lower primary grades emphasize perception and experience, upper primary and middle school emphasize understanding and application, and high school emphasizes project creation and frontier applications" [2]. This framework can be read as a structured re-expression of that official guidance on a two-dimensional grid of capability strand by grade band.

The framework rests on three design principles:

1. **Spiraling, not repeating**: The same theme (say, "data") is revisited at different levels of abstraction across grade bands, avoiding the trap of "explaining what AI is from scratch every year." This aligns with the top-level design of "tiered, spiraling progression" in the General AI Education Guidelines [3].
2. **Concepts first, tools as a replaceable vehicle**: Curricula anchor on transferable concepts and modes of thinking, treating specific tools or models as interchangeable carriers — this guards against rapid technical obsolescence eating into course content.
3. **Ethics and governance run throughout**: Ethics and governance literacy is not an "add-on unit" reserved for older grades — it runs, age-appropriately, from lower primary onward, developing in step with capability-building. The General AI Education Guidelines list "safety and controllability" as one of four basic principles, providing direct policy support for this design choice [3].

4.2.1 Primary School: Embodied Experience and Demystification

The core task at the primary level is building an accurate intuitive sense of "intelligence" — steering clear of two opposite misconceptions: deifying AI as all-knowing, or dismissing it as "just looking things up." The General AI Education Guidelines position the primary stage as "building experience and interest" [3], and Beijing's Work Plan for Advancing AI Education in Primary and Secondary Schools (2025–2027) goes further, specifying that the primary stage should be "primarily experiential coursework, focused on introduction and awakening interest" [7][8]. Teaching should center on activities that are embodied, observable, and comparative: showing a model's boundaries by letting it get things wrong, understanding the division of labor between humans and machines by comparing how each completes the same task, and building intuition about the relationship between data and patterns through activities like "teaching the machine to recognize things."

It is worth stressing that at this stage, students' independent, unsupervised use of open-ended generative tools should be tightly restricted. This is not merely a conservative pedagogical preference — it is an explicit policy requirement. The Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition) stipulate that at the primary level, "students are prohibited from independently using open-ended content-generation functions," and that teachers may only "make appropriate use of it as a classroom teaching aid" [4]. Accordingly, "using AI" at the primary level should take place within closed scenarios designed by teachers, while students' own independent exploration should focus on underlying intuitions about data, patterns, and rules — rather than directly driving generative models themselves.

4.2.2 Middle School: Understanding Mechanisms and Establishing Norms

Middle school carries the transition from "knowing how to use it" to "understanding a bit of why it works," and it is also the critical period for establishing academic-integrity norms in earnest. The General AI Education Guidelines position this stage around "understanding the logic of the technology," "engineering thinking," and "deepened ethical awareness" [3]; Beijing's plan frames it as "primarily awareness-building coursework, focused on guiding students to use AI to empower their own learning and daily life" [7]. On content, this stage introduces the basic mechanics of machine learning — the intuitions behind classification, prediction, and training versus inference — alongside its inherent limitations (dependence on data, potential for bias, lack of genuine semantic understanding). On practice, it moves students from "casually asking" to "using AI in a structured way," introducing foundational prompt-engineering methods and explicit design of human–AI division of labor.

On usage boundaries, the Guidelines for Generative AI Use permit middle-school students to "moderately explore logical analysis of generated content" [4] — meaning students may, under teacher guidance, treat generated output as "material to be critically analyzed" rather than "an answer to be copied." This dovetails with this stage's goal of understanding mechanisms and recognizing limitations. This is also the stage at which workable academic-integrity rules need to be established (see Section 4.5).

4.2.3 High School: Systems-Level Understanding and Agent Orchestration

High school targets systems-level understanding and complex problem-solving. The General AI Education Guidelines position this stage as "technology strategy," "innovative application," "systems thinking," and "practicing social responsibility" [3]; Beijing's plan frames it as "primarily comprehensive and practice-oriented coursework, focused on strengthening students' AI application ability and innovative spirit" [7]. The Guidelines for Generative AI Use permit high-school students to "conduct inquiry-based learning grounded in technical principles" [4] — that is, more open-ended inquiry, built on an understanding of how models actually work. Content at this stage reaches into the working mechanisms of generative models and agents, along with the sources of their capabilities and risks; practice moves into orchestrating agents across multi-step tasks, comprehensive cross-disciplinary applications, and engineering implementation. At this stage, students should be guided to move from a "user" perspective to an "evaluator/designer" perspective — able to critically verify AI system outputs and produce structured assessments of their social impact, deeply coupled with governance literacy.

4.3 Prompt Engineering and Agent Orchestration as Curriculum Content

The spread of generative AI and agents has brought two new content areas into the K-12 curriculum's field of view: **prompt engineering** and **agent orchestration**. Turning either into coursework risks two pitfalls: narrowing prompt engineering into "phrasing tricks" or a "grab-bag of magic incantations," and discussing method in the abstract, detached from actual subject-matter content.

This Blue Book argues both should be framed as **"modes of thinking for human–AI collaboration"** rather than standalone skills:

- **The pedagogical core of prompt engineering** is task decomposition, clarifying intent, expressing constraints, and verifying results — all transferable higher-order thinking skills that connect directly to subject-area writing, inquiry, and problem-solving. Courses should carry this content through authentic subject-matter tasks rather than through a dedicated "prompting class." This approach aligns with the General AI Education Guidelines' emphasis on "solving real-world problems" within the skills dimension, and with the 2024 Notice's advocacy for "task-based, project-based, and problem-based learning" [2][3].
- **The pedagogical core of agent orchestration** is breaking a complex goal into delegable subtasks, designing the human–AI division of labor and handoff points, and maintaining oversight and correction over the automated process. This overlaps substantially with project management, systems thinking, and computational thinking, and is a natural extension of the "systems thinking" objective for high school [3].

One useful reference point is a capability scale proposed by industry: the "AI Application Capability" tiers for education. When NetEase Youdao launched a new suite of products in its Ziyue education-LLM series in August 2025, it introduced an L1–L5 capability-tier standard for AI application in education, describing the field as currently "accelerating from L3 active-learning tutoring toward L4 virtual teacher" [9]. This tiering scheme comes from a vendor context and has not yet been independently vetted by researchers, but it usefully signals something for curriculum designers: an agent's "degree of automation" sits on a continuous spectrum, and students need to learn to judge how much decision-making authority to hand to an agent, and at which handoff points a human must take over. That judgment — not proficiency with any single product — is the core literacy that makes agent orchestration worth teaching.

In this domain, first-person classroom observation captured by teachers and analysis of human–AI collaboration process data provide useful support for "treating process as an object of instruction." Relevant methods are discussed in the corresponding chapters of this Institute's [TBD: name of relevant AI glasses in education report, 2026] and [TBD: name of relevant education robotics report, 2026].

4.4 Comparing K-12 AI Curriculum Policy Across China, the US, and the EU

Curriculum rebuilding does not happen in a vacuum — it is shaped by each country's (or region's) curriculum standards, policy direction, and governance framework. This section presents a structured comparison of K-12 AI curriculum policy across three representative regions: China, the United States, and the European Union. All data below is drawn from official primary-source documents; where a claim cannot be traced to a primary source, it is marked [TBD].

Dimension	China	United States	European Union
Governing policy/document	Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) and Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition), preceded by Document No. 32 [2024] of the General Office of Basic Education and the "AI+Education" Action Plan under Document No. 1 [2026] of the Department of Science, Technology and Informatization [1][3][4][5]	Executive Order 14277, "Advancing Artificial Intelligence Education for American Youth" (signed 2025-04-23) [10][11]	The EC/OECD joint AI Literacy Framework (AILit), "Empowering Learners for the Age of AI" (finalized 2026-06-18), and the EC's Guidelines on the Ethical Use of AI and Data in Education [12][13]
Governance-tier characteristics	Centrally coordinated: ministry-level guidelines plus locally developed curriculum guides (e.g., Beijing's plan), "specifying course objectives, content, and class-hour requirements for each grade band" [5][7]	Federal steering with state/district autonomy: the executive order establishes a White House AI Education Task Force and uses grants, public–private partnerships, and competitions as levers, while curriculum authority remains with states and districts [10][11]	Member-state primary responsibility with an EU-level shared framework: AILit is a non-mandatory reference framework that countries may voluntarily incorporate into their own curricula and policies [12]
Curriculum positioning	General education, which may take the form of "standalone courses, cross-disciplinary integration, or	No unified federal required course; policy goals center on "advancing AI literacy, teacher training, and early	A literacy-framework-level reference, with actual implementation left to

	hands-on activities" in parallel [3][7]	exposure" [10][11]	member states [12]
Literacy/curriculum framework	An integrated "four-in-one" literacy spanning knowledge, skills, thinking, and values [3]	The executive order calls for "AI literacy and proficiency" without specifying a unified competency list [10]	Four domains — "Engage with," "Create with," "Manage," and "Shape AI" — spanning knowledge, skills, and attitudes [12][13]
Stance on generative-AI governance	A dedicated Guidelines for Generative AI Use sets usage boundaries by grade band (prohibited for independent use in primary school; moderate exploration in middle school; inquiry-based use in high school) [4]	The executive order emphasizes "promoting adoption and literacy"; generative-AI governance is handled mainly through state/district policy and existing privacy law (e.g., FERPA) [10]	EC ethical guidelines require teachers to use AI and data responsibly, and fold generative AI's impact on disinformation into digital-literacy guidance [13]
Teacher preparation/training mechanisms	AI-education teacher training is built into teacher-training programs, with education bases established in phases and labs allocated on an equitable basis [1][2]	The executive order directs the Department of Education to prioritize AI within competitive teacher-training grants, and the NSF to prioritize funding for AI-education research [10][11]	The EC has issued multiple guideline sets for teachers; per OECD figures, roughly 40% of teachers had received AI-related training as of 2024 [13][14]
Adoption-scale data	Essentially universal AI education in primary and secondary schools targeted before 2030; Beijing requires no fewer than 8 class-hours per academic year starting fall 2025, spanning primary through high school [1][7][8]	Khanmigo (a Microsoft partnership, free for US teachers) reached roughly 2 million users worldwide in SY24-25, with registrations up 731% year over year, including 770,000 K-12 students through US school-district partnership programs [15]	Public consultation ahead of the AILit framework's finalization drew input from 100+ countries and 2,000+ participants; the framework is intended for adoption into national curricula and policy [12]

The comparative analysis yields the following conclusions:

· **Governance-path differences are empirically stark**: The three regions occupy clearly distinct positions on the spectrum from centralized coordination to local autonomy. China follows a top-down path of "ministry-level guidelines setting the standard, with local curriculum guides specifying class-hours" — Beijing's "no fewer than 8 class-hours per year, spanning all grade bands" is a concrete local-implementation example [5][7]. The United States follows a decentralized path in which the federal government steers through executive orders, grants, and competitions while curriculum authority sits with states and districts [10]. The EU follows a collaborative path in which member states hold primary responsibility while the EU supplies a shared reference framework (AILit) [12]. This divergence means "best practice" cannot simply be transplanted across borders — it must be localized to fit each region's governance structure.

· **A shared pivot from "teaching AI" to "governing AI"**: A common thread across all three regions' recent policy is the addition of governance concerns — "how to use AI responsibly in education" — alongside the pre-existing question of "how to teach AI." China addresses this through its dedicated, grade-banded Guidelines for Generative AI Use [4]; the EU through ethical guidelines requiring responsible use [13]; the United States relies more heavily on existing privacy law and state-level policy [10] — a reflection of the practical pressure created once generative AI actually reached classrooms.

· **Curriculum standards lag behind the technology**: Formal curriculum-standard revision cycles typically run on a scale of years, while model capabilities iterate on a scale of months. The EU's AILit is a case in point: its draft was published in May 2025 and finalized in June 2026, a span of roughly a year [12] — during which domestic and international education LLMs went through several iterations of their own. This "curriculum-versus-technology" time lag effectively pushes part of the burden of keeping pace onto localities, schools, and individual teachers.

Note: This section's cross-national comparison relies on official primary-source documents. The claim that the AILit framework comprises "19 competencies" appears in media coverage [14]; the official finalized version specifies only that it is organized into "four domains, structured around knowledge, skills, and attitudes," without giving a running total [12]. The main text accordingly follows the official framing of "four domains" and makes no quantitative claim about the total number of competencies. Where a claim cannot be traced to a primary source, it is marked TBD rather than estimated.

4.5 Building Academic-Integrity and AI-Use Norms into the Curriculum

Generative AI has put pressure on the traditional assumption of "independent completion" underlying academic work, turning academic integrity from a matter of after-the-fact discipline into one of up-front design. Building academic-integrity norms into the curriculum is one of the clearest markers distinguishing an AI-native curriculum from a tool-embedded one. Policy already provides a concrete handle for this shift: the Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition) require students to "avoid simply copying content generated by generative AI tools into assignments" and to "avoid using generative AI in exams and tests, and must not use generative AI to cheat," while also drawing clear lines for teachers — "generative AI must not be used as a substitute for the teacher as the primary instructional agent," and teachers should "avoid using AI-generated content directly to evaluate students" [4].

This section proposes a **tiered usage framework** as a practical handle for making explicit how much AI a given learning task permits. The L0–L3 tiers below are this Blue Book's own operational framework; their boundaries are designed to be compatible with the grade-banded provisions of the Guidelines for Generative AI Use — in particular, the primary-level prohibition on "independent student use of open-ended content-generation functions" corresponds to a stricter default tier [4]:

Usage Tier	Degree Permitted	Typical Applicable Tasks	Disclosure Requirement
L0 – Prohibited	No AI use whatsoever	Closed-book exams, tests, basic skills assessments (corresponds to the Guidelines' "must not use generative AI to cheat" / "avoid use in exams and tests") [4]	Not applicable
L1 – Restricted	Limited to designated supporting steps	Reference lookup, terminology clarification, and similar supporting steps, but not drafting core content (corresponds to "avoid simply copying generated content") [4]	Must declare which steps used AI
L2 – Collaborative	Collaboration permitted, with substantive human involvement required	Co-drafting followed by substantive revision and critical verification by the student (echoes middle school's "moderate exploration of logical analysis of generated content") [4]	Must document the human–AI division of labor
L3 – Open	Full use encouraged	Exploring tool capabilities; inquiry-based learning into underlying principles (echoes high school's "inquiry-based learning grounded in technical principles") [4]	Must explain how AI was used

Supporting elements for building this into the curriculum include:

· **Disclosure over prohibition**: Rather than imposing sweeping bans that are difficult to monitor, schools should build clear, low-cost disclosure mechanisms, shifting the emphasis from "catching cheaters" to "clarifying the rules and assessing genuine ability." This aligns with the Guidelines' own governance stance of "guarding against over-reliance on AI tools" and "preventing the digital divide caused by technology dependence" [4].

· **Data and privacy red lines alongside integrity norms**: Academic-integrity norms cannot be separated from data security. The Guidelines explicitly state that it is "strictly prohibited to input personal information, exam questions, or other sensitive data into AI tools" and instruct users to "avoid inputting personal information into generative AI tools" [4] — so any disclosure mechanism should cover three layers at once: whether AI was used, how it was used, and what sensitive information (if any) was entered.

· **Coordinated reform of assessment methods**: Academic-integrity norms must evolve in step with assessment methods (see Chapter 5, "Assessment in an AI Environment," of this Blue Book) — for example, by adding more process-based evidence, oral defenses, and in-person generation tasks that AI cannot easily complete on a student's behalf.

· **Age-appropriate calibration of norms**: Integrity education cannot rely on institutional rules alone — it must be calibrated to cognitive development at each grade band, moving from an introduction to "honest use" in primary school up to norms around "attribution, citation, and responsibility" in high school — a progression structurally consistent with the grade-banded loosening built into the Guidelines [4].

4.6 Education-Vertical LLMs: Model Supply as Curriculum Infrastructure

Because AI-native curricula default to models and agents as infrastructure, "which model to use" has risen from a technical-selection question to a curriculum-and-governance question. This judgment is now backed by national policy: the "AI+Education" Action Plan explicitly calls for "the state to organize coordinated research efforts, developing AI education LLMs by education stage, and strengthening capabilities such as value alignment, logical reasoning, and safety ethics" [5][6]. Relative to general-purpose LLMs, **education-vertical LLMs** — models adapted through education-specific data, alignment, and productization — carry specific expectations in the K-12 setting.

4.6.1 What Sets Education-Vertical Models Apart

Model supply aimed at K-12 settings carries differentiated requirements that diverge from general consumer or workplace use cases, chiefly in the following respects:

· **Safety and age-appropriateness**: Content safety, value alignment, and protection of minors are hard constraints that take priority over raw capability ceilings. The Action Plan lists "value alignment" and "safety ethics" among the core capabilities targeted for education-LLM research, lending direct policy backing to this priority [5].

· **Pedagogical alignment**: Model behavior needs to match instructional intent — for example, Socratic-style guidance that prompts rather than hands over the answer directly, which sits in tension with general-purpose assistants oriented toward "getting to the answer in one step." International practice already offers a point of comparison here: Google has built LearnLM directly into Gemini, reporting that expert reviewers preferred LearnLM over GPT-4o roughly 31% more often on average across a range of learning-scenario evaluations — with "alignment for learning" as its central selling point [16][17].

- **Controllability and auditability**: Teachers need the ability to intervene in model behavior, and student usage needs to be traceable to support assessment and governance. The Guidelines explicitly prohibit teachers from "using AI to answer student questions directly" and prohibit inputting sensitive data into tools — requirements that, in practice, demand that model supply take a controllable, auditable product form [4].
- **Data compliance and localization**: Because education data involves minors' privacy and compliance obligations, data-handling practices and deployment models face higher requirements, often pointing toward private or on-premises deployment capability.

4.6.2 On-Device Deployment and Delivery Models

Consistent with the broader on-device trend discussed throughout this Blue Book, deployment of education-vertical models is evolving from a single cloud-based model toward a coordinated "cloud–edge–device" structure. On-device/localized deployment offers distinct value for K-12 settings — keeping data on campus, enabling low-latency interaction, supporting offline availability, and controlling cost — but device compute constraints force trade-offs between model scale and capability. A workable path is typically a layered supply of "cloud-based large models plus on-device small/distilled models," routed by task sensitivity and real-time requirements.

Industry has already produced product forms that echo this trend: educational smart hardware — AI learning tablets, answer-assistance pens, and similar devices — has become an important vehicle for delivering education LLMs. For example, NetEase Youdao's "Youdao AI Answer Pen Space X," launched in August 2025, packages the Ziyue LLM's capabilities into a handheld device; the company's official figures state that it handles "more than 10 question-answering interactions per user per day on average" [9]. Zuoyebang and Yuanfudao (Xiaoyuan) have likewise connected their learning tools and hardware to reasoning models such as DeepSeek, blending them with their own proprietary models [18]. This kind of "model-on-device" practice is a concrete expression of on-device intelligence in the education setting. This trend corroborates the account of on-device intelligent hardware's classroom applications given in this Institute's [TBD: name of relevant AI glasses in education report, 2026].

4.6.3 Overview of the Education-Vertical Model Supply Landscape

Providers currently supplying vertical models and agent-based teaching assistants for education fall broadly into three categories: education editions from general-purpose LLM vendors, proprietary or adapted models from education-technology companies, and localized adaptations of open-source or third-party foundation models. The table below lists representative providers whose information is publicly verifiable; specific market-share figures and other quantitative data remain subject to official disclosure and are marked TBD where not independently verified — no list entries have been fabricated.

Supply-Side Category	**Typical Positioning**	**Representative Providers (publicly verified)**	**Adoption/Share**
General vendors' education editions	General capability plus education alignment/guardrails	Google LearnLM (built into Gemini, aligned for learning) [16][17]; Khanmigo (Khan Academy, Microsoft partnership, free for US teachers) [15][19]	Khanmigo: roughly 2 million users worldwide in SY24-25, 770,000 K-12 students through US school-district programs, registrations up 731% year over year, coverage in 70+ countries [15]; remainder [TBD]
Ed-tech proprietary/adapted models	Deep scenario focus plus platform integration	iFlytek Spark (education edition) [20]; TAL Education's "Jiuzhang" LLM [21]; NetEase Youdao's "Ziyue" [9][22]; Zuoyebang's "Yinhe" LLM, Yuanfudao's (Yuanli) LLM, and others (all have completed LLM registration/filing) [18][22]	[TBD: specific adoption/share figures pending official disclosure]
Open-source/third-party foundation-model localization	Controllable, low-cost, privately deployable	Education companies connecting to reasoning models such as DeepSeek and blending them with proprietary models (e.g., Zuoyebang, Xiaoyuan) [18]	[TBD]

Note: In the table above, TAL Education's "Jiuzhang" LLM has been widely reported as "the first math-focused LLM" and described as performing strongly on math problem-solving benchmarks; this reflects vendor and media framing rather than independent performance verification, and is recorded here purely as a factual note on the supply landscape [21]. NetEase Youdao has proposed an L1–L5 capability tiering for

AI applications in education [9]. Any claims involving superlatives, market share, or performance rankings are attributed to their original vendor/media source; this Blue Book does not offer an independent rating.

4.6.4 Evaluation Dimensions for Selection and Governance

To avoid having curricula locked into a single vendor or swept along by an immature product, schools and districts should build a structured evaluation framework before adopting an education-vertical model as curriculum infrastructure — covering at minimum:

1. **Pedagogical fit**: Whether the model supports "guided, Socratic-style prompting" rather than "answers handed over directly," and whether it works well with task-based, project-based, and problem-based learning (echoing the instructional orientation of the 2024 Notice [2]).
2. **Content safety and age-appropriateness**: Whether content filtering, value alignment, and protection of minors meet required standards (echoing the Action Plan's emphasis on "value alignment, safety ethics" [5]).
3. **Data compliance and privacy protection**: Whether the model meets requirements that data stay on campus and that sensitive information not leak out (echoing the data red lines in the Guidelines [4]).
4. **Explainability and interventability**: Whether teachers can trace, audit, and correct model behavior.
5. **Cost and sustainability**: Including total cost of ownership for on-device or private deployment.
6. **Vendor stability and portability**: Avoiding vendor lock-in while preserving the ability to migrate across platforms.

This evaluation framework connects to the procurement and accountability mechanisms discussed in Chapter 6, "Governance." It is worth noting that education LLMs remain in a period of rapid iteration and intense commercial competition — product names, capability boundaries, and market structure can all shift within short cycles. Selection decisions should therefore treat portability as a first-order constraint from the outset, not as a remedy applied after the fact.

4.7 Chapter Summary

This chapter defined an AI-native curriculum as one organized around literacy as its central axis, the human–AI collaborative process as its core object of instruction, and education-vertical models as its default infrastructure, systematically rebuilt across grade bands. It proposed a two-dimensional grade-band-by-capability rebuilding framework, discussed the pedagogical core of prompt engineering and agent orchestration as curriculum content, presented an evidence-based comparison of curriculum policy across China, the United States, and the EU, set out a tiered framework for academic-integrity use norms, and — from the supply side — analyzed the differentiated requirements, on-device deployment trends, and selection-and-governance considerations for education-vertical LLMs as curriculum infrastructure. Several points deserve emphasis:

- **The substance of curriculum rebuilding is a change in how questions are asked**, not simply an increase in class-hours or tools layered on top of existing content. China's 2024–2026 policy progression — from general-education guidelines, to usage guidelines, to organized research into education LLMs — embodies exactly this three-part coordination: from "teaching AI," to "establishing norms for using AI," to "building purpose-built education models" [1][3][4][5].
- **Grade-band progression should spiral, lead with concepts, and carry ethics throughout**, guarding against technical obsolescence while fitting how students actually develop cognitively. This orientation runs in the same direction as the General AI Education Guidelines' "tiered, spiraling progression" and "four-in-one literacy," and as the EU AILit framework's "four domains" [3][12].
- **Model supply has become a curriculum variable in its own right** — "what to teach" and "what to teach it with" now have to be designed and governed together. The dense rollout of domestic and international education-vertical models (Spark, Jiuzhang, Ziyue, Yinhe, and others, alongside LearnLM and Khanmigo), together with the national-level mandate to "develop education LLMs by education stage," has elevated this variable from an industry talking point to a matter of curriculum and governance [5][9][15][16][21].

Several specifics referenced throughout this chapter — particular market shares, product-capability rankings, local class-hour details, and the total number of competencies in some international frameworks — remain marked [TBD] wherever they could not be traced to a primary source; no figures or proper nouns have been fabricated. Building on this chapter, Chapter 5 turns to rebuilding learning assessment for an AI environment, and Chapter 6 turns to the governance mechanisms that must accompany curriculum and model supply.

Chapter References

[1] Notice of the General Office of the Ministry of Education on Strengthening Artificial Intelligence Education in Primary and Secondary Schools (Document No. 32 [2024] of the General Office of Basic Education) [教育部办公厅关于加强中小学人工智能教育的通知（教基厅函〔2024〕32号）]. Ministry of Education of the People's Republic of China. 2024.

http://www.moe.gov.cn/jyb_xwfb/gzdt_gzdt/s5987/202412/t20241202_1165500.html

[2] Ministry of Education Announces: Essentially Universal AI Education in Primary and Secondary Schools Before 2030 [教育部发文：2030 年前中小学基本普及人工智能教育]. China Education

and Research Network (CERNET). 2024.
https://www.edu.cn/xxh/focus/zc/202412/t20241202_2644322.shtml

[3] Officially Released! Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [正式发布！《中小学人工智能通识教育指南（2025 年版）》来了]. China Education and Research Network (CERNET). 2025.
https://www.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[4] Here's the Full Text: Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition) [速看！《中小学生成式人工智能使用指南（2025 年版）》全文来了]. China Education and Research Network (CERNET). 2025.
https://www.edu.cn/xxh/focus/zc/202505/t20250513_2667992.shtml

[5] Notice of the Ministry of Education and Four Other Departments on Issuing the "AI+Education" Action Plan (Document No. 1 [2026] of the Department of Science, Technology and Informatization) [教育部等五部门关于印发《"人工智能+教育"行动计划》的通知（教科信〔2026〕1 号）]. Ministry of Education of the People's Republic of China. 2026.
http://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html

[6] Notice of the Ministry of Education and Four Other Departments on Issuing the "AI+Education" Action Plan [教育部等五部门关于印发《"人工智能+教育"行动计划》的通知]. China Higher Education Student Information (CHSI). 2026.
https://www.chsi.com.cn/jyzx/202604/20260410/2293459694.html

[7] Notice of the Beijing Municipal Education Commission on Issuing the Beijing Work Plan for Advancing AI Education in Primary and Secondary Schools (2025–2027) [北京市教育委员会关于印发《北京市推进中小学人工智能教育工作方案（2025—2027 年）》的通知]. Beijing Municipal Education Commission. 2025.
https://jw.beijing.gov.cn/xxgk/2024zcwj/2024qtwj/202503/t20250307_4028227.html

[8] Primary and Secondary Schools to Offer AI General-Education Courses Starting Fall Semester, No Fewer Than 8 Class-Hours Per Year, Covering All Grade Bands from Primary Through High School [中小学秋季学期起开设 AI 通识课 每学年不少于 8 课时 覆盖小学至高中全学段]. Beijing Municipal Government Portal (Capital Window). 2025.
https://www.beijing.gov.cn/ywdt/gzdt/202503/t20250308_4029252.html

[9] NetEase Youdao Launches Multiple New Ziyue Education-LLM AI Products, Defines L1–L5 Tiers for AI Application Capability in Education [网易有道发布子曰教育大模型多款 AI 新品 定

义教育 AI 应用能力 L1-L5 分级]. Xinhua News Agency. 2025.

http://www.news.cn/tech/20250821/11dbed39f03b4b3aa0723e0a2ce910d9/c.html

[10] Executive Order 14277—Advancing Artificial Intelligence Education for American Youth. The White House. 2025. https://www.whitehouse.gov/presidential-actions/2025/04/advancing-artificial-intelligence-education-for-american-youth/

[11] Executive Order 14277—Advancing Artificial Intelligence Education for American Youth (signed 2025-04-23). The American Presidency Project, UC Santa Barbara. 2025. https://www.presidency.ucsb.edu/documents/executive-order-14277-advancing-artificial-intelligence-education-for-american-youth

[12] New AI Literacy Framework Helps Schools Prepare Learners for the Age of Artificial Intelligence / Empowering Learners for the Age of AI (AILit Framework, finalized 2026-06-18). European Commission (European Education Area). 2026. https://education.ec.europa.eu/whats-new/news/new-ai-literacy-framework-helps-schools-prepare-learners-for-the-age-of-artificial-intelligence

[13] Guidelines on the Ethical Use of Artificial Intelligence and Data in Teaching and Learning. European Commission (European Education Area). 2024–2025. https://education.ec.europa.eu/focus-topics/digital-education/actions/plan/ethical-guidelines-for-educators-on-using-artificial-intelligence

[14] European Commission and OECD Set 19 AI Literacy Competences for Schools. EdTech Innovation Hub. 2026. (Media report; source of the "19 competencies" figure — the official finalized version does not itemize a running total.) https://www.edtechinnovationhub.com/news/european-commission-and-oecd-set-19-ai-literacy-competences-for-schools

[15] Khan Academy Annual Report SY24-25 (Khanmigo usage scale: roughly 2 million users; 770,000 K-12 students through US school-district programs; registrations up 731% year over year; coverage in 70+ countries). Khan Academy. 2025. https://annualreport.khanacademy.org/

[16] I/O 2025: LearnLM in Gemini 2.5 and More AI Updates to Help People Learn (LearnLM built into Gemini; roughly 31% expert preference rate over GPT-4o). Google (The Keyword blog). 2025. https://blog.google/products-and-platforms/products/education/google-gemini-learnlm-update/

[17] LearnLM: Improving Gemini for Learning. Google / arXiv:2412.16429. 2024. https://arxiv.org/abs/2412.16429

[18] Yuanfudao's Full Product Line Connects to DeepSeek, with Tuning Results to Roll Out Over the Following Two Months [小猿全系产品接入 DeepSeek，未来两个月内陆续推出调优成果].

Duozhi.com. 2025. http://www.duozhi.com/industry/insight/2025021116990.shtml

[19] Microsoft, Khan Academy Provide Free AI Assistant for All US Teachers (Microsoft partnership, free for US teachers). CNBC. 2024. https://www.cnbc.com/2024/05/21/microsoft-khan-academy-launch-free-ai-assistant-for-all-us-teachers.html

[20] iFlytek Spark LLM: Latest Progress and Typical Applications in Education [讯飞星火大模型最新进展及其在教育领域典型应用]. Chinese Association of Automation. 2024. https://www.caa.org.cn/article/345/5225.html

[21] From Large Models to Robots: Companies Intensify Their "AI+Education" Push (covers TAL Education's Jiuzhang LLM, Seewo's teaching LLM, and other supply-side entrants) [从大模型到机器人 企业密集布局"AI+教育"（好未来九章大模型、视源希沃教学大模型等供给格局）]. Xinhua News Agency. 2025. http://www.news.cn/tech/20250609/d168be922c8449e0bf91b94a97c67403/c.html

[22] The Age of Personalized Learning Has Arrived! Yuanfudao, Xueersi, and Other Education LLMs Complete Regulatory Filing [个性化学习时代来了！猿辅导、学而思等教育大模型通过备案]. Southern Metropolis Daily (Nandu). 2024. https://m.mp.oeeee.com/a/BAAFRD000020240605961163.html

Chapter 5 Agentic Teaching Assistants and the Pedagogy of Human–AI Co-Learning

5.1 From Tool-Assisted Instruction to Agentic Collaboration: Locating the Paradigm Shift

Back in 2022, when the master version of this Blue Book was first conceived, generative AI had not yet entered mainstream K-12 classrooms. AI in education was understood mainly as algorithmic recommendation, automated grading, and adaptive practice — **closed, one-directional tools**. The relationship between student and system was one of "use and be used": the system responded to inputs according to preset rules, neither generating open-ended content nor initiating interaction on its own.

Since 2023, the maturation of large language models (LLMs) and multimodal foundation models — compounded by the accelerating **agentic** turn of 2025–2026 — has qualitatively changed what education technology can be. The agentic teaching assistant is no longer a passive question-answering bot. It is a software entity capable of **understanding goals, decomposing tasks, invoking tools, sustaining multi-turn memory, and acting autonomously**. Under goals set by a teacher, it can continuously probe, diagnose, demonstrate, and give feedback to an individual student or small group, approaching the interactional density of one-on-one tutoring.

Authoritative frameworks have already captured this shift. The U.S. Department of Education's Office of Educational Technology, in its May 2023 report *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations*, held that AI should **augment rather than replace human intelligence** and placed "keeping humans in the loop" first among its seven core recommendations — teachers, not AI, remain accountable for instructional decisions, assessment, and student support; AI is a tool, not an authority [1]. UNESCO's September 2023 *Guidance for Generative AI in Education and Research* — the first global guidance of its kind — went further, offering countries a human-centered policy framework built on agency, inclusion, equity, and linguistic and cultural diversity, and, unprecedentedly, recommending a **minimum age of 13** for classroom use of AI tools alongside a call for accompanying teacher training [2]. Together, these two documents establish the baseline for any discussion of agentic teaching assistants: leaps in technical capability must be bounded by human-centered governance.

The shift, then, is not about replacing teachers — it is about reconfiguring the three-way relationship of **human–machine–human** in teaching. This chapter's core claim is that the value of introducing agentic teaching assistants at the K-12 level depends heavily on the maturity of the accompanying

Human-AI Co-learning Pedagogy. An agent deployed without a pedagogical framework tends to slide into a shortcut for outsourcing homework; an agent embedded in well-designed pedagogy, by contrast, can amplify a teacher's professional judgment and extend the reach of personalized instruction. Mutlu Cukurova of University College London offered a theoretical anchor for this distinction in a 2024 *British Journal of Educational Technology* article introducing the lens of "Hybrid Intelligence." He distinguishes three roles AI can play in education — externalizing human cognition, internalizing an AI model to shape human thinking, and **extending cognition** through a genuinely hybrid human–AI system — and argues that only the third path expands human cognition while preserving teachers' professional judgment and students' agency [3].

> On how edge-side agents are being realized in hardware form factors, see this institute's *Global Blue Book on AI Glasses in Education 2026*; on the intersection of embodied agents and educational robotics, see this institute's *Global Blue Book on Educational Robotics Development 2026*.

This paradigm shift has found clear uptake in Chinese policy as well. In November 2024, the General Office of China's Ministry of Education issued the *Notice on Strengthening AI Education in Primary and Secondary Schools* (Document No. 32 [2024] of the General Office of Basic Education [教基厅函〔2024〕32 号]), setting an overarching goal of **essentially universalizing AI education in primary and secondary schools by 2030** and laying out six tasks, including building a systematic curriculum framework, normalizing instruction and assessment, and scaling up teacher supply [4]. In April 2026, five ministries — the Ministry of Education, the National Development and Reform Commission, the Ministry of Industry and Information Technology, the Ministry of Science and Technology, and the National Data Administration — jointly issued the *"AI+Education" Action Plan* (Document No. 1 [2026] of the Ministry of Education's Department of Science, Technology and Informatization [教科信〔2026〕1 号]), which explicitly calls for developing an "intelligent learning companion," exploring "human-AI collaborative teaching models," advancing "intelligent grading, Q&A, and tutoring," and planning a national AI-computing service platform for education alongside education-stage AI foundation models [5]. Human–AI co-learning, in other words, has moved from an academic concept to a stated direction of national education-digitalization strategy.

5.2 A Capability-Function-Role Framework for Agentic Teaching Assistants

To avoid discussing agentic assistants in the abstract, this section proposes a three-layer **capability–function–role** framework for characterizing where assistants of different maturity levels fit in the classroom.

5.2.1 Capability Tiers

Along the technical-capability dimension, agentic teaching assistants can be roughly divided into four progressive tiers:

Tier	Capability signature	Typical instructional interaction	Intensity of teacher involvement	Real-world product/system examples
L1 Responsive	Single-turn Q&A, knowledge retrieval	Student asks, gets an instant explanation	Low	Generic chatbot Q&A interfaces
L2 Tutoring	Multi-turn dialogue, Socratic follow-up questioning, error diagnosis	Guided problem-solving, concept clarification	Medium	Khanmigo student-facing tool [6]; Jill Watson forum Q&A [7]
L3 Task-oriented	Task decomposition, tool invocation, process memory	Project-based learning companionship, staged scaffolding	Medium-high	LearnLM triggered-intervention tutoring [8]; agentic teaching assistants
L4 Collaborative	Long-term cross-session memory, learner profiling, proactive intervention	Personalized learning pathways, early-warning tutoring	High (requires governance)	Mostly research prototypes/pilots; large-scale evidence still lacking

It bears emphasizing that the higher the tier, the more stringent the requirements around **data retention, algorithmic transparency, and educational ethics**. Publicly available products currently aimed at K-12 mostly deliver real capability at L1–L2. Even the more evidence-backed examples — Khan Academy's Khanmigo and Google DeepMind's LearnLM, developed in partnership with Eedi — perform reliably in real classrooms mainly at multi-turn guided tutoring (L2) and a controlled, constrained form of triggered intervention (L3), not autonomous long-term intervention across sessions (L4) [6][8]. Any K-12 product claiming full L4 autonomous collaboration still needs its evidence scrutinized carefully; this Blue Book does not classify any specific product as a verified L4 case.

5.2.2 Mapping Capability to Instructional Function

Mapped onto concrete instructional moments, the functions an agentic teaching assistant can plausibly take on include, but are not limited to:

- **Conceptual scaffolding**: providing graduated hints ("hint laddering") rather than direct answers when a student's cognitive load runs high. A randomized controlled trial of Stanford's Tutor CoPilot found that tutors given AI assistance were **more likely to use guiding questions and**

less likely to hand over direct answers — a behavioral shift consistent with high-quality teaching practice [9];

- **Formative feedback**: giving immediate, specific, actionable feedback on open-ended responses to shorten the feedback loop. In LearnLM's UK classroom trial, **74.4%** of the tutoring messages it generated were approved as-is by human supervising tutors, with roughly another 2% needing only minor edits; the factual-error rate was just 0.1% (5 of 3,617 messages), and harmful content was zero [8];
- **Differentiated practice generation**: dynamically producing problems and scenarios calibrated to a student's current level;
- **Metacognitive prompting**: prompting students to restate their reasoning, question themselves, and evaluate their own solutions, building learning strategies;
- **Reducing teacher workload**: assisting with lesson preparation, generating tiered instructional materials, and summarizing learning data. Khanmigo offers teachers **more than 25** dedicated tools (generating lesson ideas, personalizing assignments, building grading rubrics, summarizing student progress, and more) and, since August 2024, has been **free** for U.S. K-12 teachers with support from Microsoft Azure OpenAI Service (previously a $4/month fee), later extending to more than 180 countries [10].

Empirical effect sizes for these functions are accumulating quickly, but **results vary enormously by subject, grade band, and fidelity of implementation** — and not always in the same direction. The same technology, configured well, can meaningfully improve learning; configured poorly, it can just as easily undermine it (see the evidence review in Section 5.4). No single effect size, then, should be cited as if it generalized beyond the conditions under which it was produced.

5.3 Four Design Principles for Human–AI Co-Learning Pedagogy

The educational value of agentic teaching assistants does not materialize on its own — it depends on **deliberate pedagogical design**. This section distills four design principles for K-12 settings into an operational framework for schools and district instructional-research offices. These are not abstract normative aspirations; each is grounded directly in primary evidence from the past two years.

5.3.1 Principle One: Don't Outsource Cognition (the Boundaries of Cognitive Offloading)

The primary risk in human–AI co-learning is that students hand off **core cognitive labor** that they should be doing themselves to the agent, resulting in a hollowed-out competence that looks like learning but isn't. This is no longer a theoretical worry — it is an empirical phenomenon that large-sample studies have repeatedly confirmed.

The most representative case is a large-scale randomized controlled trial by Hamsa Bastani, Osbert Bastani, and colleagues at the University of Pennsylvania, conducted at a Turkish high school and published in the *Proceedings of the National Academy of Sciences* (PNAS) in June 2025. Nearly 1,000 students in grades 9–11 used one of two GPT-4-based tutoring tools during math practice. **While they had AI access**, a plain ChatGPT interface ("GPT Base") raised scores by 48%, and a version with pedagogical guardrails ("GPT Tutor") raised them by 127%. **But once AI access was removed**, students who had used the unguarded version actually scored **17% lower** than a control group that had never used AI at all — they had leaned on AI as a crutch during practice, and their independent competence suffered as a result. Crucially, the guardrailed version essentially eliminated this negative effect [11]. The paper's own title — "Generative AI without guardrails can harm learning" — is about as direct a citation as exists for the boundary this principle is meant to protect.

Broader evidence comes from a quasi-experimental analysis (2025) of **ten years and 3.2 million** learning interactions on the ALEKS platform. Exploiting the contrast between "AI-susceptible" tasks (word problems) and "AI-resistant" tasks (graphing problems), the study found that after ChatGPT's release, college students' study time on AI-susceptible problems fell by 2.8% per quarter, accumulating to a 26.9% decline over 11 quarters — high school students fell by 31.3%. The authors call this "cognitive surrender": students adopt AI output with minimal scrutiny, ceding cognitive control and bypassing the active learning essential to durable knowledge [12].

Design, then, should ensure AI carries **assistive cognition** (information retrieval, formatting, first-draft generation) while leaving **generative cognition** — argumentation, reasoning, creation, judgment — to the student. Workable mechanisms include requiring students to critically revise AI output, requiring them to show their thinking rather than just their results, isolating AI use during key assessments, and **building "don't just give the answer" pedagogical guardrails directly into tutoring agents** — Bastani and colleagues' research shows that this single layer of prompt design is exactly what determines whether AI helps or harms learning [11].

5.3.2 Principle Two: Teacher-in-the-Loop

Agentic teaching assistants should be positioned as an **extension of the teacher, not a replacement** — squarely in line with the U.S. Department of Education's 2023 recommendation to "keep humans in the loop" [1]. Pedagogical design needs to guarantee teachers' **visibility, control, and veto power** over agent behavior: teachers should be able to review student–agent interaction logs, adjust the agent's guidance strategy, and step in promptly to correct inappropriate or erroneous content.

The most compelling evidentiary case for "teacher-in-the-loop" is Stanford's Tutor CoPilot study, led by Rose Wang and colleagues — **the first randomized controlled trial of a human–AI system in real online tutoring**, covering 900 tutors and 1,800 K-12 students from historically underserved communities. Here, AI does not face the student directly; instead, it acts as an **"expert co-pilot"**

embedded in real time in the human tutor's workflow, suggesting high-quality teaching language to the tutor during sessions. The result: tutors with access to Tutor CoPilot saw their students' probability of mastering a topic rise by **4 percentage points** ($p<0.01$). Tellingly, the students who benefited most were those with **lower-rated tutors** — a relative mastery gain of **9 percentage points** over the control group — and the system cost only about **$20 per tutor per year** [9]. This case demonstrates that injecting AI capability **into human teacher judgment**, rather than routing around teachers to act on students directly, can both amplify professional practice and preserve accountability — arguably the most scalable realization of "teacher-in-the-loop" available today.

A deployment without a teacher-in-the-loop mechanism, in essence, shifts instructional responsibility onto a system that cannot be held accountable. By contrast, LearnLM's UK classroom trial used a "human-supervised, AI-drafted" architecture, in which 17 supervising tutors retained the right to approve, edit, or reject every AI-generated message. In the end, 82.4% of tutors said the system's greatest value was its ability to "support multiple students simultaneously," and their comfort rating for the AI tool rose from 3.4 to 3.9 (out of 5) — an example of "teacher-in-the-loop" delivering gains in both efficiency and trust at once [8].

5.3.3 Principle Three: Explainability and Contestability

Agents designed for minors should, at the pedagogical level, cultivate in students the habit of mind of **treating AI as an interlocutor to be questioned, not taken at face value**. This requires the agent to surface its own uncertainty at appropriate moments, cite its sources, and acknowledge the limits of its competence — and it requires course design that has students actively verify AI output and catch and correct AI errors (treating "AI getting it wrong" as an instructional resource in its own right).

The case for this principle rests on quantitative grounds. Even LearnLM — pedagogically optimized and human-supervised — still produced factual errors in 0.1% of its messages [8]; Jill Watson's accuracy on a synthetic test set ranged from **75% to 97%** depending on the content source [7]. Both findings point to the same conclusion: "AI makes mistakes" is a baseline condition that classrooms must address explicitly, not an edge case that can be waved away. Building the verification of AI output and the identification of AI errors into learning tasks themselves serves double duty — it is part of building AI literacy, and it hedges pedagogically against the risk of model hallucination (the AI-literacy assessment framework is detailed in Chapter 3 of this Blue Book). China's *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)* likewise lists "human–AI collaboration competency" as a core literacy, emphasizing at the high school level an understanding of generative AI's social impact and limitations — providing a curricular entry point for cultivating this contestability mindset [13].

5.3.4 Principle Four: Equity by Design

Human–AI co-learning implemented without attention to equity risks widening rather than narrowing the **digital divide**: disparities in devices, connectivity, family support, and teachers' own AI literacy could let already-advantaged students benefit disproportionately.

Notably, the latest evidence suggests equity can cut **either way**. On one hand, rolled out without guardrails, AI is more likely to be used as a "crutch" by students with weaker self-regulation and a greater tendency to reach for shortcuts, widening the gap [11][12]. On the other, well-designed AI can generate a disproportionately larger positive effect for **disadvantaged groups**: in the Tutor CoPilot study, it was precisely the students of **lower-rated tutors** (many serving under-resourced communities) who gained the most (+9 percentage points) [9]. This suggests human–AI collaboration can function as a mechanism for **redistributing capability** — delivering high-quality teaching practice, at low cost, to where it is needed most. That finding reframes "equity of access" from a passive risk to guard against into an active pedagogical design target.

Pedagogical design should explicitly account for accessibility across urban–rural divides, students with special educational needs, and linguistic-minority communities, including offline/edge-device deployment options, accessible interaction design, and low-barrier pathways to use. Chinese policy is already pushing in this direction: the *Notice on Strengthening AI Education in Primary and Secondary Schools* lists "urban–rural coordination" as one of four guarantee measures [4]; Microsoft's partnership with Khan Academy, through donated cloud computing, has made Khanmigo's teacher-facing tools free in more than 180 countries and discounted for economically disadvantaged schools — a commercial-track realization of "equity of access" [10]. (A systematic analysis of equity of access appears in the relevant dedicated chapter of this Blue Book.)

5.4 Three Classroom Models for Human–AI Co-Learning

Building on the principles above, this section summarizes three classroom organizational models that have gradually emerged in practice between 2024 and 2026, offered as reference points for schools designing instructional activities. All three already correspond to real products or pilots, but **adoption scale and effect data must be read strictly against the evidence, and results are highly dependent on implementation fidelity**.

Model One: AI-as-Tutor (Personalized Tutoring)

Here the agent acts as a one-on-one tutor for the student, providing differentiated guidance under goals the teacher sets — well suited to practice consolidation and targeted remediation or enrichment. This is the model with the strongest evidence base to date. Khan Academy's Khanmigo is the leading example: as of late 2024, roughly **266 U.S. school districts** had deployed it, student usage jumped from 40,000 to 700,000 over the 2024–25 school year and was projected to surpass a million in

2025–26, and districts such as Newark, New Jersey have piloted it at roughly $35 per student per year [6]. A crossover randomized controlled trial by Harvard's Gregory Kestin and colleagues, published in *Scientific Reports* in June 2025, offers the strongest effect-size evidence yet for this model: 194 college physics students using a custom AI tutor ("PS2 Pal," built on GPT-4) achieved learning gains **more than double** those from active-learning classroom instruction, with an effect size of roughly **0.73–1.3 standard deviations**, and **in less time** (a median of 49 minutes for the AI group versus 60 minutes in class, with 70% of students finishing within an hour) — alongside higher engagement and motivation scores [14]. **Its boundaries need to be flagged carefully**, however: the sample was Harvard undergraduates, the study ran for only two weeks, and it focused on mid-level cognitive skills such as understanding, application, and analysis — it cannot speak to long-term retention, transfer, or synthetic creativity [14], let alone be extrapolated directly to K-12. The model's success still hinges on avoiding degeneration into an answer-dispensing machine, with Socratic guidance and graduated hints as its core mechanism — precisely the value of the "guardrails" that Bastani and colleagues demonstrated [11].

Model Two: AI-as-Peer (Collaborative Inquiry)

Here the agent acts as a "virtual peer" within a learning group — joining discussion, raising counterexamples, and taking on different positions to spark argumentation and critical thinking. It suits project-based learning (PBL) and open-ended inquiry. Georgia Tech's Jill Watson is the long-running real-course example of this model: deployed as a virtual teaching assistant embedded in course discussion forums, it ran across five courses (including "Introduction to Programming") in spring 2024, with Q&A accuracy significantly better than general-purpose assistants. It also **improved students' sense of "teaching presence" and "social presence,"** and classes using it saw a modestly higher proportion of A grades (66% vs. 62%) along with a small uptick in academic performance [7]. This case suggests that the value of a "virtual peer/assistant" often extends well beyond answering questions correctly — it can also **improve the social-psychological climate of the classroom**. This model depends especially on Principle Three (contestability), since students tend to accept a peer-presenting AI's claims less critically.

Model Three: AI-as-Co-creator (Collaborative Creation)

Here students and the agent jointly complete writing, coding, design, and other creative tasks, with AI handling drafts, iteration, and technical scaffolding while students own the concept, decisions, and final sign-off. This model is tightly coupled to academic-integrity concerns and carries **the highest risk of cognitive offloading** — "cognitive surrender" is most likely to occur in open-ended creative tasks [12]. It therefore requires clear norms for disclosing and attributing AI use, alongside design that requires students to leave a visible trail of their thinking and to substantively revise AI output (academic-integrity norms are detailed in Chapter 4 of this Blue Book).

These three models are not mutually exclusive; a single classroom can combine them depending on the instructional stage. It is worth stressing that most of the effect-size evidence cited above **comes from higher education or short-duration trials** — large-scale, long-duration, cross-disciplinary controlled studies at the K-12 level remain thin. LearnLM's UK classroom RCT (165 students aged 13–15, 7 weeks) is one of the few high-quality studies aimed directly at the secondary-school age range, and its conclusion is that "AI tutoring can safely and effectively support students": the LearnLM group matched or slightly outperformed human tutoring on immediate remediation (93.0%), misconception resolution (95.4%), and, notably, **knowledge transfer** (66.2%, 5.5 percentage points above human tutoring) [8]. Even so, the researchers themselves characterize the study as "exploratory," a reminder that K-12-level evidence is still in its early stages of accumulation.

5.5 Implementation Risks and the Interface with Governance

The pedagogical deployment of agentic teaching assistants inevitably touches on risks that require a governance-level response. This section offers only a pedagogy–governance interface note; the systematic governance framework appears in this Blue Book's governance chapter.

· **Learning dependency and skill atrophy**: the risk with the most solid evidence base — unguarded AI use drives a 17% drop in independent scores once access is removed [11], and large-sample data show a "cognitive surrender"-style contraction of study time [12]. This should be constrained through assessment design (isolating AI at key checkpoints), boundaries on use, and built-in pedagogical guardrail prompts;

· **Bias and content safety**: models can output biased, age-inappropriate, or factually incorrect content. Even under human supervision, LearnLM still produced factual errors in 0.1% of messages [8], so content filtering, human review, and age-appropriateness mechanisms remain indispensable; UNESCO's recommended minimum age of 13 can serve as a policy reference point for age-appropriateness [2];

· **Data privacy**: the collection, retention, and use of interaction data from minors must follow principles of minimal necessity and informed parental consent. UNESCO's guidance calls on countries to adopt data-protection and privacy standards [2]; China's *"AI+Education" Action Plan* proposes building digital student profiles and dynamically optimizing learning pathways, which in turn means compliance governance over minors' learning data must be strengthened in step [5] (this chapter does not speculate about the specific domestic policy document number governing minors' data protection where applicable; that verification is left to the dedicated governance chapter);

· **Algorithmic transparency and accountability**: proactive interventions by L3–L4 agents must be traceable, explainable, and appealable, consistent with the U.S. Department of Education's "humans in the loop" recommendation [1];

· **Teacher professional development**: teachers need new competencies for directing agents and designing human–AI co-learning activities. Evidence shows teacher comfort with AI can improve substantially through practice (rising from 3.4 to 3.9 in the LearnLM trial) [8], and the *Notice on Strengthening AI Education in Primary and Secondary Schools* already lists "scaling up teacher supply" as a core task [4] (teacher AI-literacy development is detailed in this Blue Book's dedicated teacher chapter).

5.6 Summary

This chapter examines agentic teaching assistants through the lens of **Human–AI Co-learning Pedagogy**. Its central claim is that an agent's educational value does not derive from its technical capability alone, but from the pedagogical design and governance built around it. The 2024–2026 primary evidence reviewed here confirms this claim in something close to a controlled-experiment pattern: the same GPT-4 tutoring, with pedagogical guardrails, improved learning — without them, it undermined learning [11]; the same AI intervention, when it injected itself into teacher judgment (Tutor CoPilot), delivered the largest benefit to disadvantaged groups [9] — when it bypassed the teacher, it tended toward cognitive surrender [12]. Through its capability–function framework, four design principles, and three classroom models, this chapter offers a structured set of analytical tools for "responsibly introducing agentic teaching assistants" at the K-12 level, anchored in real products (Khanmigo, Jill Watson, Tutor CoPilot, LearnLM) and authoritative studies (Harvard's *Scientific Reports*, PNAS, the Stanford RCT). Subsequent chapters build further on the dimensions of assessment (Chapter 6) and governance (Chapter 7), completing a full literacy–curriculum–human/AI collaboration–assessment–governance chain of capability and governance. Every specific adoption figure, effect size, and policy document number cited in this chapter has been traced to a real source; where a corresponding domestic policy document number could not be verified, the text uses qualitative language rather than any fabricated number.

Chapter References

[1] Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations · U.S. Department of Education, Office of Educational Technology · 2023 · https://www.ed.gov/sites/ed/files/documents/ai-report/ai-report.pdf

[2] Guidance for Generative AI in Education and Research (includes the recommendation of a minimum age of 13 for classroom use) · UNESCO · 2023 · https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research ; related coverage: https://www.unesco.org/en/articles/unesco-governments-must-quickly-regulate-generative-ai-schools

[3] Cukurova, M. "The interplay of learning, analytics and artificial intelligence in education: A vision for hybrid intelligence." · *British Journal of Educational Technology*, 56 · 2024 · https://www.scirp.org/reference/referencespapers?referenceid=3952753

[4] Notice of the General Office of the Ministry of Education on Strengthening AI Education in Primary and Secondary Schools (Document No. 32 [2024] of the General Office of Basic Education, proposing essential universalization by 2030; Chinese: 教育部办公厅关于加强中小学人工智能教育的通知，教基厅函〔2024〕32 号) · Ministry of Education of the People's Republic of China · 2024 · http://www.moe.gov.cn/jyb_xwfb/gzdt_gzdt/s5987/202412/t20241202_1165500.html

[5] Notice of Five Ministries Including the Ministry of Education on Issuing the "AI+Education" Action Plan (Document No. 1 [2026] of the Department of Science, Technology and Informatization; Chinese: 教育部等五部门关于印发《"人工智能+教育"行动计划》的通知，教科信〔2026〕1 号) · Ministry of Education, National Development and Reform Commission, Ministry of Industry and Information Technology, Ministry of Science and Technology, National Data Administration · 2026 · http://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html

[6] Khanmigo/Khan Academy AI tutoring: 266 districts, student usage 40,000→700,000, $35/student/year pilot · Chalkbeat Newark (2024) and CBS 60 Minutes (2024-12) · 2024 · https://www.chalkbeat.org/newark/2024/05/13/artificial-intelligence-khanmigo-chatbot-tutor-pilot-testing-districtwide-expansion/ ; https://www.cbsnews.com/news/khanmigo-ai-powered-tutor-teaching-assistant-tested-at-schools-60-minutes-transcript/

[7] Kakar, S. et al. "Jill Watson: A Virtual Teaching Assistant Powered by ChatGPT" (deployed across five courses in spring 2024, accuracy 75–97%, improved teaching/social presence, A-grade rate 66% vs. 62%) · Georgia Tech / AI-ALOE · 2024 · https://tech.ai.gatech.edu/news/georgia-techs-jill-watson-outperforms-chatgpt-real-classrooms ; paper: https://dl.acm.org/doi/10.1007/978-3-031-64302-6_23

[8] LearnLM Team (Google DeepMind & Eedi). "AI tutoring can safely and effectively support students: An exploratory RCT in UK classrooms" (N=165, ages 13–15, 5 schools, 7 weeks; approval rate 74.4%, factual-error rate 0.1%, knowledge transfer 5.5 percentage points above human tutoring) · arXiv:2512.23633 · 2025 · https://arxiv.org/html/2512.23633v1

[9] Wang, R. E. et al. "Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise" (first real-tutoring human–AI RCT, 900 tutors/1,800 K-12 students; +4 p.p., lower-rated tutors +9 p.p., $20/tutor/year) · Stanford National Student Support Accelerator / arXiv:2410.03017 · 2024 · https://nssa.stanford.edu/studies/tutor-copilot-human-ai-approach-scaling-real-time-expertise ; https://arxiv.org/abs/2410.03017

[10] Khanmigo for Teachers (25+ teacher tools, free in the U.S. since August 2024 with Azure OpenAI support, previously $4/month, expanded to 180+ countries) · Microsoft Education Blog / Khan Academy Blog · 2024 · https://www.microsoft.com/en-us/education/blog/2024/08/khanmigo-for-teachers-your-free-ai-powered-teaching-tool/ ; https://news.microsoft.com/source/features/ai/khan-academy-and-microsoft-partner-to-expand-access-to-ai-tools/

[11] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., Mariman, R. "Generative AI without guardrails can harm learning: Evidence from high school mathematics" (RCT with nearly 1,000 grade 9–11 students in Turkey; +48%/+127% with AI access, −17% for the unguarded group after removal) · *PNAS* 122 · 2025 (working paper July 2024) · https://www.pnas.org/doi/10.1073/pnas.2422633122 ; working paper: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4895486

[12] Rismanchian, S., Uzun, H., Matayoshi, J., Cosyn, E., Kurd-Misto, E. "Faster Completion, Less Learning: Generative AI Reduced Study Time on Math Problems and the Knowledge They Build" (ten years and 3.2 million interactions on ALEKS; study time on AI-susceptible problems −2.8% per quarter, cumulative −26.9% for college students and −31.3% for high school students; "cognitive surrender") · arXiv:2605.21629 · 2026 · https://arxiv.org/abs/2605.21629

[13] Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) (Ministry of Education Basic Education Teaching Guidance Committee; includes "human–AI collaboration competency" as a core literacy, alongside both standalone courses and cross-disciplinary integration; Chinese: 《中小学人工智能通识教育指南（2025 年版）》) · Ministry of Education / CERNET · 2025 · https://www.cernet.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[14] Kestin, G. et al. "AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting" (194 Harvard physics students, crossover RCT, PS2 Pal built on GPT-4; learning gains more than double active learning, effect size approximately 0.73–1.3 SD, less time required) · *Scientific Reports* · 2025 · https://www.nature.com/articles/s41598-025-97652-6 ; commentary and effect-size detail: https://etcjournal.com/2025/11/10/review-of-kestin-et-al-s-june-2025-harvard-study-on-ai-tutoring/

Chapter 6 Prompt-Engineering Literacy, Academic Integrity, and Campus Norms for Generative AI Use

6.1 Introduction: From "Can Use the Tool" to "Uses It Well, Uses It Wisely"

Now that generative AI has entered K-12 classrooms, the central question facing the enabling infrastructure is no longer whether students can access AI tools, but whether students and teachers can use AI responsibly, productively, and verifiably. The 2022 master edition of this white paper was written before generative AI saw mass adoption, so its literacy framework centered on understanding the principles behind algorithms and data. The reality of 2024–2026 is different: conversational generative models and agentic assistants have become tools that students reach for as a matter of course. A Pew Research Center survey of American teens aged 13–17, fielded in fall 2024, found that about **26%** said they had used ChatGPT to help with schoolwork — double the 13% share recorded in 2023. Yet teens are far from agreed on which uses are acceptable: 54% considered it acceptable to use AI for "researching new topics," while only 18% considered it acceptable to use it for "writing essays," and 42% said that use was flatly unacceptable[1]. These figures capture exactly what this chapter is concerned with: AI use has become routine, but social consensus on what counts as legitimate use is still far from settled.

This shift has pulled three governance domains — once handled separately — into tight coupling: **prompt-engineering literacy** (interacting with models effectively and critically), **academic integrity** (where "reasonable assistance" ends and "academic misconduct" begins), and **campus use norms** (how schools institutionalize the regulation of AI use). This chapter treats the three as one mutually reinforcing governance system: literacy supplies the underlying capability, integrity supplies the value floor, and norms supply the institutional guarantee. Leave any one out, and AI use on campus tends to become inefficient, unmoored from norms, or simply chaotic.

This "literacy–integrity–norms" coupling is already visible in international policy texts. UNESCO's *Guidance for generative AI in education and research*, published in September 2023 as the world's first such guidance, calls on governments to regulate generative AI "swiftly" through seven key steps, among them safeguarding data privacy, setting a minimum age (recommended no lower than 13) for independent use of generative AI platforms, and training teachers[2]. The same guidance treats capability-building, ethical boundaries, and institutional constraints as inseparable — a genuine international precedent for the three-part coupling this chapter builds on.

For how generative AI is being applied within specific subject teaching and how agentic teaching assistants are evolving technically, see Chapters 3 (Human–AI Co-Teaching) and 4 (Assessment and Evaluation) of this institute's *Global Blue Book on the Enabling Infrastructure for AI Education in K-12 Schools 2026*. This chapter focuses on the literacy–integrity–norms governance dimension.

6.2 Prompt-Engineering Literacy as an Emerging Foundational Literacy

6.2.1 Defining Prompt-Engineering Literacy

In a K-12 context, prompt engineering should not be reduced to "memorizing a set of prompt templates" or "mastering a knack for phrasing." Recent scholarship has firmly repositioned it as a composite foundational literacy rather than a single operational skill. Federiakin and colleagues, writing in *Frontiers in Education* (2024), frame "prompt engineering as a new 21st century skill" — the ability to state a problem, its context, and the constraints on the desired answer clearly enough to elicit a fast, accurate response from an AI assistant. They break this down into four interlocking dimensions: (1) understanding basic prompt structure (instruction, context, input data, output indicator); (2) **prompt literacy** — precision of phrasing and avoidance of pitfalls such as ambiguity and insufficient context; (3) prompting methods (techniques such as few-shot, chain-of-thought, and tree-of-thought prompting); and (4) **critical online reasoning** — evaluating model output against established standards[3]. The framework singles out iteration as prompt engineering's defining feature: users must "toggle repeatedly between issuing a request and evaluating the output" to revise or abandon a given line of inquiry, and critical evaluation of that output depends on domain knowledge as a "critical benchmark," with keeping a human in the loop remaining essential throughout[3].

Drawing on this research and on K-12 practice, this chapter redefines prompt-engineering literacy as a set of transferable cognitive and practical capabilities:

- **Intent clarification** — turning a vague learning need into a structured, concrete task description;
- **Context provisioning** — judging what background information, constraints, and examples the model needs to complete the task, and supplying them;
- **Iterative refinement** — following up and revising purposefully in response to shortcomings in the model's output, forming a "ask–evaluate–ask again" loop;
- **Critical verification** — treating model output with appropriate skepticism, spotting hallucinations, factual errors, and logical gaps, and corroborating externally;
- **Boundary awareness** — understanding the model's capability limits and appropriate use cases, and knowing when AI should not be used at all.

Of these, **critical verification and boundary awareness** are what separate prompt-engineering literacy from mere tool operation. The former points toward responsible adoption of AI output and is the cognitive precondition for academic integrity. A systematic review focused on higher education (*International Journal of Educational Technology in Higher Education*, 2025) reaches a consistent conclusion after surveying the prompt-engineering literature: prompt engineering is not just about getting better outputs — it is a vehicle for building critical thinking, evaluation, and reflective capacity, and should be designed into curricula rather than treated as an incidental skill[4].

6.2.2 How Prompt-Engineering Literacy Relates to Existing Literacy Frameworks

Prompt-engineering literacy is not a freestanding new subject; it belongs inside existing AI literacy and digital literacy frameworks. It overlaps substantially with critical thinking, information literacy, and media literacy, and can be understood as these established literacies made concrete in a generative-AI setting.

Among the more influential framings of AI literacy as an overarching construct: Long and Magerko (2020) define AI literacy as "a set of competencies that enables individuals to critically evaluate AI technologies, communicate and collaborate effectively with AI, and use AI as a tool online, at home, and in the workplace"[5]. Ng and colleagues (2021) connect AI literacy to Bloom's taxonomy, conceptualizing it across four dimensions: knowing and understanding AI, using and applying AI, evaluating and creating AI, and AI ethics[6]. Prompt-engineering literacy sits squarely across three of these — using/applying, evaluating/creating, and ethics — combining the capacity to use AI with the disposition to evaluate its output and use it responsibly. Embedding prompt-engineering literacy within existing AI-literacy frameworks, rather than treating it as a standalone subject, is therefore the more economical and logically coherent path.

Related literacy dimension	Overlap with prompt-engineering literacy	Corresponding overarching framework
Information literacy	Verifying output, tracing sources, cross-checking	Ng et al. (2021), "evaluating and creating AI" dimension[6]
Critical thinking	Spotting hallucinations, evaluating logic, detecting bias	Federiakin et al. (2024), "critical online reasoning" dimension[3]
Digital citizenship	Privacy protection, data boundaries, accountability	Ng et al. (2021), "AI ethics" dimension[6]; Long & Magerko (2020)[5]
Subject-matter literacy	Domain knowledge as the basis for judging output correctness	Federiakin et al. (2024), "domain knowledge as critical benchmark"[3]

One important mechanism to flag: **prompt-engineering literacy depends heavily on subject-matter knowledge as its verification base**. A student who lacks basic understanding of a subject has no way to judge whether the model's output is right or wrong, and the prompt–verify loop collapses. This mechanism is directly supported by the research cited above — critical evaluation requires domain knowledge as a "critical benchmark"[3]. The implication is that building AI literacy cannot substitute for solid subject-matter education; it must be built on top of it. The two are complementary, not interchangeable.

6.2.3 A Grade-Band Progression for Capability Building

Given how much cognitive development varies across age groups, building prompt-engineering literacy should follow a staged, grade-band progression. This approach lines up closely with national curricula in major education systems. Take China as an example: the *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)*, issued in May 2025 by the Ministry of Education's Committee on Basic Education Instruction Guidance, lays out a staged AI-literacy framework — primary school "emphasizes cultivating interest and building basic awareness," focusing on perceptual exposure to technologies such as speech recognition and image classification; lower-secondary school "strengthens understanding of technical principles and basic applications," covering the basic machine-learning workflow and the concept of supervised learning, and explicitly requires students to "identify the risk of misinformation in generative AI applications"; upper-secondary school "emphasizes systems thinking and innovative practice," building understanding of the technical characteristics and social impact of generative AI[7]. Notably, the critical-verification skill of "identifying misinformation" is placed at the lower-secondary stage, while "understanding the characteristics and impact of generative AI" is placed at the upper-secondary stage — a progression from awareness to critique.

Drawing on international experience, this chapter proposes the following qualitative grade-band progression; specific hours, assessment indicators, and materials should be determined separately in line with each country's curriculum:

- **Primary school**: centered on "awareness and safety" — building the basic understanding that "AI can be wrong and needs an adult to help check it," with an emphasis on safe boundaries of use. This corresponds to the "cultivating interest and basic awareness" orientation in China's curriculum guidelines[7].
- **Lower secondary**: centered on "method and verification" — mastering structured questioning and basic output-verification methods, understanding the basic requirements of academic integrity, and being able to "identify the risk of misinformation"[7].
- **Upper secondary**: centered on "critique and governance" — understanding how models work and where their limits lie, engaging with data and bias issues, integrating AI responsibly into authentic tasks, and understanding the technical characteristics and social impact of generative AI[7].

It is also worth noting that making AI literacy a statutory obligation for the general population is becoming a regional trend. Article 4 of the EU AI Act, in force since 2 February 2025, requires providers and deployers of AI systems to ensure that their staff and others dealing with the operation of AI systems on their behalf have "a sufficient level of AI literacy," calibrated to their technical knowledge, experience, education, and the context of use[8]. Although this provision binds institutions rather than school curricula directly, it marks a shift of "AI literacy" from an educational aspiration to a legally binding social competency — and provides useful institutional backdrop for literacy-building at the K-12 level.

6.3 Academic Integrity: Redefined for the Generative-AI Era

6.3.1 Why the Traditional Academic-Integrity Framework Strains Under Generative AI

Generative AI poses a structural challenge to the traditional academic-integrity framework. That framework is built around "plagiarism" — using someone else's existing expression without attribution — and its determination depends on there being a comparable original source to check against. Generative models, by contrast, produce text synthesized on the fly, typically without a single traceable source, which means the source-comparison logic underlying traditional detection tends to break down.

The resulting core tension is this: reasonable AI use (brainstorming, language polishing, clarifying questions) and academic misconduct (having AI write the work, fabricating results) are not separated by a natural bright line — they sit on a **continuous spectrum**. The same AI-use behavior can carry a different integrity valence depending on the learning objective and assessment context. A systematic survey of generative-AI and large-language-model guidelines at 80 universities worldwide (published in *Nature Human Behaviour*, 2025) confirms how widespread this tension is: institutions are shown repeatedly weighing an optimistic case — fostering collaborative creativity, expanding educational access, empowering teachers and students — against concerns about ethical complexity, balancing innovation against academic integrity, unequal access, and the risk of misinformation, reflecting the absence of any unified consensus on where the boundaries of AI use should lie[9].

6.3.2 A "Use-Spectrum" Analytical Framework

To give schools an analytical tool for policymaking, this chapter proposes a "use spectrum" framework for generative AI, organized around **learning objectives** and **use transparency**, that classifies AI-use behaviors by their impact on learning goals (the table below is an analytical framework; each school should adapt it to its own assessment objectives):

Use type	Typical behavior	Impact on learning objective	General integrity determination
Cognitive-support	Concept explanation, prompting inquiry,	Supports understanding; the learner remains the	Generally acceptable; disclosure advised

	answering questions	cognitive agent	
Production-support	Language polishing, structural suggestions, translation/proofreading	Neutral to positive, depending on what is being assessed	Depends on the task's objective; disclosure required
Partial-substitution	Generating a draft, doing part of the reasoning for the student	May hollow out the very competency being assessed	High risk; requires explicit authorization and disclosure
Full-substitution	Having AI write the work, answer for the student, or fabricate process and data	Fully hollows out learning and assessment	Academic misconduct

The central claim of this framework is that **integrity determinations should be tied to the specific learning objective being assessed, rather than applying a blanket ban or blanket permission to AI use**. When the assessment target is precisely the core competency the student is meant to acquire, using AI to substitute for that competency constitutes misconduct; when AI use does not touch the core competency being assessed, it can be permitted under a transparent disclosure requirement. This "objective-bound" approach is corroborated by the operating rules of major examination bodies. The UK's Joint Council for Qualifications (JCQ), in its guidance *AI Use in Assessments: Your role in protecting the integrity of qualifications*, states plainly that for non-examined assessments, students using AI as a source of information **must record the name of the AI tool used and the date the content was generated, retain a non-editable copy of the questions and the AI-generated content (such as a screenshot), and submit a brief description of how it was used alongside the work**; failing to declare AI use — for a student who has already been instructed on referencing conventions and has signed an authenticity declaration — can result in a penalty ranging from a mark deduction to a mark of zero, and potentially more severe sanctions including disqualification[10]. This shows that whether the assessed competency is touched, and whether use is transparently disclosed, together determine the integrity status of the very same behavior.

6.3.3 Disclosure as a Governance Lever

Given how unreliable detection methods currently are, **disclosure** is emerging internationally as the key governance lever for AI integrity in education. The underlying logic: rather than pouring resources into increasingly unreliable after-the-fact detection, shift governance upstream to prior agreement and in-the-moment disclosure, requiring students to proactively and honestly reveal how and to what extent they used AI.

Typical elements of a disclosure mechanism include whether AI was used at all, which tool was used, at which stages, and retention of the key prompts and outputs. This mechanism has already been written into several major education systems' official frameworks. Beyond the JCQ requirement noted above — naming the tool, recording the generation date, and retaining a non-editable copy[10]

— Australia's *Australian Framework for Generative Artificial Intelligence in Schools*, approved by state and territory education ministers in October 2023, establishes "transparency" as one of its core principles: where the use of generative AI tools will affect others, that use should be disclosed — to teachers, staff, students, parents, or carers as appropriate — and the framework strongly encourages schools to proactively identify and manage academic-integrity risks, including by setting expectations, building a culture of integrity, factoring generative AI into assessment design, and identifying and responding to inappropriate use[11]. It should be noted that a disclosure mechanism's effectiveness depends on a supporting culture of integrity and sound assessment design; a bare disclosure requirement, without a foundation of trust and educational guidance, risks becoming a mere formality.

6.3.4 The Limits of AI-Detection Tools, and a Case for Caution

This chapter takes a cautious stance on technical detection of "AI-generated text." Existing AI-text detectors have inherent accuracy limitations, and **false positives** — misclassifying human writing as AI-generated — can unfairly harm students, with a documented systematic bias against non-native writers. A study by Liang and colleagues at Stanford University (*Patterns*, 2023) provides direct evidence: when tested against TOEFL essays, seven mainstream GPT detectors misclassified **61.3%** of essays written by non-native English speakers as AI-generated; **97.8%** of these essays were flagged by at least one detector, and **19.8%** were flagged by all seven — yet every one of these essays was written entirely by a human. By contrast, the same detectors' false-positive rate on essays written by native English-speaking eighth-graders in the US was close to zero[12]. The study's mechanistic explanation is that detectors judge largely on surface features such as text "perplexity," and non-native writers tend to use simpler, more regular vocabulary and sentence patterns — which makes their writing more likely to be flagged as machine-generated. In effect, the detectors "cannot distinguish 'writing in a second language' from 'being generated by a machine'"[12]. Given that students in China and across much of the non-English-speaking world are precisely the population most exposed to this systematic bias, this evidence carries direct policy relevance for this white paper.

This chapter therefore recommends: output from AI-detection tools should be treated, at most, as a **supplementary reference**, and should never serve as the sole or decisive evidence for a finding of academic misconduct. Integrity governance should instead lean more heavily on improved assessment design — process-oriented assessment, oral defense, and key steps completed in class — rather than on an arms race in technical detection.

6.4 Campus Norms for Generative AI Use: Elements of Institutional Design

6.4.1 Basic Principles for Drafting Norms

School-level norms for generative AI use should avoid two extremes: unrestricted permissiveness on one side, and a simplistic blanket ban on the other. This "avoid both extremes" stance is already explicit in national-level policy. China's *Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition)*, issued in 2025 by the Ministry of Education's Committee on Basic Education Instruction Guidance, is built on the principle of "application first, governance as the foundation," focusing on application scenarios such as instructional support, personalized learning, and intelligent administration, while emphasizing data security and ethical bottom lines and proposing a "government–school–enterprise–family–society" multi-party coordination mechanism[7][13]. The UK Department for Education's (DfE) policy paper *Generative artificial intelligence (AI) in education*, updated in June 2025, takes a similarly "conditionally encouraging" stance: it emphasizes that, used safely and effectively with appropriate infrastructure in place, AI can reduce teachers' administrative burden and support feedback and personalized learning — while stating plainly that the evidence on the benefits and risks of students' own use of generative AI is still emerging, and that the DfE will keep working with the sector to clarify effective, safe use cases[14].

Building on this analysis, norm-drafting should follow these principles:

- **Education first**: the primary purpose of any norm is to promote, not obstruct, learning, with the ultimate goal of cultivating responsible users;
- **Contextual tiering**: differentiate by grade band, subject, task type, and assessment context rather than applying a single blanket rule;
- **Transparency and verifiability**: replace unreliable covert detection with a disclosure mechanism, building toward a traceable record of use (echoing both the JCQ retention requirement[10] and the Australian framework's transparency principle[11]);
- **Equity of access**: ensure the norm does not create new unfairness stemming from differences in students' household access to AI resources;
- **Dynamic revision**: given how fast the technology is moving, norms should include a mechanism for periodic review and revision.

For the equity implications of urban–rural and special-population gaps in access to AI resources, see Chapter 5 (Equity and Accessibility) of this institute's *Global Blue Book on the Enabling Infrastructure for AI Education in K-12 Schools 2026*; the norm-design principles in this chapter and the equity arguments in that chapter reinforce one another.

6.4.2 Core Elements of a Campus AI Use Norm

A workable campus norm for generative AI use should generally cover the following element modules (the specific content of each module should be filled in against local policy and law):

Element module	Core question it should answer	Example upstream reference
Scope of application	Which tools, scenarios, and people are covered	Australia's framework applies to school leaders, teachers, staff, service providers, parents, and students[11]
Tiered authorization	Which tasks permit, conditionally permit, or prohibit AI use	China's *Guidelines for Generative AI Use*, grade-band use norms[7][13]
Disclosure requirements	When and in what form AI use must be disclosed	JCQ: name the tool, record the generation date, retain a copy[10]; Australia: disclose when others are affected[11]
Data and privacy	Whether student data may be input, and how it is protected	UNESCO: adopt data-protection and privacy standards[2]; China's guidelines: data-security bottom lines[7][13]
Age and compliance	Minimum age and eligibility for each tool	UNESCO: recommends no lower than 13 for independent use; EU GDPR sets the threshold at 16[2]
Handling violations	Procedures for determining violations and educational remedies	JCQ: undisclosed use can incur penalties from a mark deduction to zero, or disqualification[10]
Teacher responsibilities	The teacher's role in guidance and oversight	EU AI Act, Article 4: deployers must ensure relevant staff have AI literacy[8]

Two of these modules — **data and privacy**, and **age and compliance** — carry legal force and must strictly align with each country's and region's minor-data-protection and cybersecurity laws; schools have no discretion here. Take the age threshold as an example: UNESCO's guidance recommends no lower than 13 for independent use of generative AI platforms, and notes that the EU's General Data Protection Regulation (GDPR) sets the consent threshold for related data processing at 16 — a difference that directly affects which tools are compliant for use in different jurisdictions[2]. When drafting norms, schools must treat the locally applicable statutory threshold as the benchmark, not a vendor's terms of service.

6.4.3 Teachers as the Linchpin of Implementation

Any campus AI-use norm ultimately depends, for its implementation, on teachers understanding it, buying into it, and enforcing it. Teachers are not merely enforcers of the norm — they also model and guide students' AI literacy and integrity awareness. This places its own demands on teachers' AI literacy: teachers need to understand the capabilities and limits of generative AI, recognize integrity risks in how students use AI, and design teaching and assessment activities that are more "AI-resilient" as a result. UNESCO's guidance lists teacher training as one of the key steps for regulating generative AI[2]. The US Department of Education's *Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations* (May 2023) is built around "humans at the center" and keeping "human in the loop" — preserving agency for educators, students, and families in educational decisions, especially high-stakes ones — emphasizing that AI should augment rather than replace teachers, and that education systems must govern their own use of AI[15]. Teacher capability is thus the hinge on which every norm's implementation ultimately turns.

> For the overall framework, training systems, and international practice around teacher AI literacy, see the dedicated Chapter 7 (Teacher AI Literacy) of this institute's *Global Blue Book on the Enabling Infrastructure for AI Education in K-12 Schools 2026*; this chapter highlights only its role within integrity governance.

6.5 Comparative Observations Across National Policy Orientations

On prompt-engineering literacy, academic integrity, and campus use norms, major education systems — China, the United States, the United Kingdom, the European Union, and Australia — show governance orientations that differ in emphasis while converging in substance. The table below structures a comparison of each system's key orientations, grounded in actual policy texts.

Dimension	China	United States	United Kingdom	European Union	Australia	International body (UNESCO)
Core policy text	*Notice on Strengthening AI Education in Primary and Secondary Schools* (Document No. 32 [2024] of the General Office of the Ministry of Education; Chinese: 教基厅函〔2024〕32 号)[16]; *Guidelines for	US Department of Education, *AI and the Future of Teaching and Learning* (May 2023)[15]; state-level guidance	DfE, *Generative AI in education* policy paper (updated June 2025)[14]; JCQ assessment guidance[10]	AI Act, Article 4, AI-literacy obligation (in force February 2025)[8]	*Australian Framework for Generative AI in Schools* (approved by education ministers, October 2023)[11]	*Guidance for generative AI in education and research* (September 2023)[2]

	General AI Education in Primary and Secondary Schools (2025 Edition)* and *Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition)*[7][13]					
Division of governance levels	Central government sets unified guidance; localities and schools implement; near-universal access targeted before 2030[16]	Federal government provides a framework; states and districts have autonomy	Central government (DfE) issues guidance; schools have autonomy	EU-level legislation binds institutions	Education ministers' council sets a unified framework; states and schools implement[11]	Global guidance; countries retain autonomy
Integrity-governance emphasis	Application first, governance as the foundation; data security and ethical bottom lines[7][13]	Human-centered; human in the loop[15]	Prevent misconduct through disclosure, retention, and assessment design (JCQ)[10]	Emphasis on institutional AI-literacy obligations[8]	Transparent disclosure plus assessment design to prevent misconduct[11]	Data privacy, age thresholds, teacher training[2]
Age/data compliance	Emphasizes data-security management norms[7]	Existing statutes such as COPPA (age 13)	Per UK data-protection law	GDPR (age-16 consent threshold)[2]	Per state privacy laws	Recommends independent use at age ≥13[2]

Three qualitative judgments follow:

· **Differences in overall orientation**: China follows an "application first, governance as the foundation" approach, emphasizing a unified national curriculum and a target of near-universal access by 2030[7][16]; the United States follows a layered pattern of federal framing with state and district autonomy; the European Union sets a legal obligation for AI literacy through legislation[8]. These systems differ structurally in how they weigh "promoting use" against "guarding against risk," and in how they divide responsibility between central authorities and local autonomy.

· **Differences in integrity-governance approach**: both the UK's JCQ and Australia's framework place the weight of governance on "transparent disclosure plus assessment design" rather than technical detection[10][11] — consistent with this chapter's cautious stance in Section 6.3.4, grounded in the evidence from Liang and colleagues (2023)[12].

· **Differences in data and age compliance**: the EU's GDPR threshold of 16 is stricter than UNESCO's recommended floor of 13[2], underscoring that the strength of legal constraints varies across systems and that schools must default to the locally applicable statutory threshold.

Notably, even in the United States, where there is no unified federal mandate, state-level governance is expanding quickly: tracking by TeachAI (an initiative launched by Code.org, ETS, ISTE, Khan Academy, the World Economic Forum, and others) together with AI for Education shows that around 26–28 states had issued K-12 AI education guidance by around 2025, a figure that had grown further to roughly 35 states plus Puerto Rico by 2026[17]. Issuing campus AI-use norms is shifting from an early-adopter behavior to a near-universal governance norm.

A common lesson that countries can draw from one another is this: **regardless of a given governance orientation's particulars, advancing literacy-building, an integrity culture, and institutional norms together is more effective than emphasizing any one of the three in isolation**. Emphasizing technical detection alone while neglecting literacy-building, or issuing norms alone without the teacher capability to back them, both tend to produce ungovernable outcomes. The *Nature Human Behaviour* analysis of guidelines at 80 universities, cited earlier, finds the same pattern: mature guidelines tend to cover capability-building, integrity constraints, and institutional procedure together, rather than favoring one at the expense of the others[9].

6.6 Chapter Summary and Recommended Direction

This chapter has treated prompt-engineering literacy, academic integrity, and campus norms for generative AI use as one mutually reinforcing governance system, yielding the following points:

1. **Prompt-engineering literacy** should be defined as a transferable foundational literacy centered on critical verification and boundary awareness. It comprises four dimensions — prompt structure, prompt literacy, prompting methods, and critical reasoning[3] — spans three dimensions of Ng and colleagues' (2021) AI-literacy framework (use, evaluation, ethics)[6], rests on subject-matter knowledge as its verification base[3], and should be built up in stages across grade bands — a progression consistent with the grade-band arrangement in China's 2025 general AI-education guidelines[7].

2. **Academic integrity**, in the generative-AI context, needs to shift from "source comparison" to "objective binding," using the "use-spectrum" framework to distinguish reasonable assistance from misconduct; disclosure serves as the governance lever (exemplified by the JCQ's tool-naming and retention requirements[10] and the disclosure principle in Australia's framework[11]); and AI-detection tools warrant a cautious stance — Stanford researcher Liang and colleagues' (2023) evidence shows a false-positive rate as high as 61.3% for non-native writers[12], which should never serve as decisive evidence of misconduct.
3. **Campus use norms** should avoid the twin extremes of permissiveness and blanket prohibition (echoing both China's "application first, governance as the foundation" stance[7][13] and the UK DfE's conditionally encouraging position[14]), following five principles — education first, contextual tiering, transparency and verifiability, equity of access, and dynamic revision — while strictly aligning with the legally binding requirements of data protection and age compliance (such as UNESCO's recommended floor of 13 and the GDPR's threshold of 16[2]).
4. **Bringing the three together**, with teacher AI literacy as the linchpin of implementation (UNESCO's teacher-training requirement[2], the AI-literacy obligation in Article 4 of the EU AI Act[8], and the US Department of Education's human-in-the-loop position[15]), is the key to sustainable governance.

For policymakers and practitioners, this chapter's recommended direction is to design the drafting of AI-use norms in tandem with literacy curricula, teacher training, and assessment reform, avoiding an approach that either "only blocks, never guides" or "only permits, never manages." Until detection methods mature into reliable tools, disclosure and improved assessment design should remain the primary levers of integrity governance. With generative AI use among adolescents already the norm — 26% of US teens reported using it for schoolwork in 2024[1] — the governance window will not stay open indefinitely. Only by establishing campus order early, through an integrated literacy–integrity–norms framework, can education hold onto its core mission of cultivating people as the technology continues to evolve.

Chapter References

[1] *About a quarter of U.S. teens have used ChatGPT for schoolwork – double the share in 2023*. Pew Research Center. 2025. https://www.pewresearch.org/short-reads/2025/01/15/about-a-quarter-of-us-teens-have-used-chatgpt-for-schoolwork-double-the-share-in-2023/

[2] *Guidance for generative AI in education and research*. UNESCO. 2023. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research (see also UN News coverage: https://news.un.org/en/story/2023/09/1140477)

[3] Federiakin, D., Molerov, D., Zlatkin-Troitschanskaia, O., & Maur, A. *Prompt engineering as a new 21st century skill*. Frontiers in Education. 2024. https://www.frontiersin.org/journals/education/articles/10.3389/feduc.2024.1366434/full

[4] *Prompt engineering in higher education: a systematic review to help inform curricula*. International Journal of Educational Technology in Higher Education (Springer). 2025. https://link.springer.com/article/10.1186/s41239-025-00503-7

[5] Long, D., & Magerko, B. *What is AI Literacy? Competencies and Design Considerations*. Proceedings of the 2020 CHI Conference. 2020. https://dl.acm.org/doi/10.1145/3313831.3376727

[6] Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. S. *Conceptualizing AI literacy: An exploratory review*. Computers and Education: Artificial Intelligence. 2021. https://www.sciencedirect.com/science/article/pii/S2666920X21000357

[7] Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [中小学人工智能通识教育指南（2025 年版）]. Committee on Basic Education Instruction Guidance, Ministry of Education. 2025. https://www.cernet.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[8] *Article 4: AI literacy, EU Artificial Intelligence Act* (applicable from 2 February 2025). European Union. 2024/2025. https://artificialintelligenceact.eu/article/4/ (see also the European Commission's AI-literacy Q&A: https://digital-strategy.ec.europa.eu/en/faqs/ai-literacy-questions-answers)

[9] Jiao, J., et al. *The global landscape of academic guidelines for generative AI and large language models*. Nature Human Behaviour. 2025. https://www.nature.com/articles/s41562-025-02124-6 (preprint: https://arxiv.org/abs/2406.18842)

[10] *AI Use in Assessments: Your role in protecting the integrity of qualifications*. Joint Council for Qualifications (JCQ), UK. 2024/2025. https://www.jcq.org.uk/exams-office/malpractice/artificial-intelligence/

[11] *Australian Framework for Generative Artificial Intelligence (AI) in Schools* (approved by the Education Ministers' Meeting, October 2023). Australian Government Department of Education. 2023. https://www.education.gov.au/schooling/announcements/australian-framework-generative-artificial-intelligence-ai-schools

[12] Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. *GPT detectors are biased against non-native English writers*. Patterns (Cell Press). 2023. https://www.sciencedirect.com/science/article/pii/S2666389923001307 (see also PMC full text: https://pmc.ncbi.nlm.nih.gov/articles/PMC10382961/)

[13] Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition) [中小学生成式人工智能使用指南（2025 年版）]. Committee on Basic Education Instruction Guidance, Ministry of Education. 2025. https://www.canedu.org.cn/site/content/7876.html

[14] *Generative artificial intelligence (AI) in education* (policy paper, updated June 2025). Department for Education (DfE), UK. 2025. https://www.gov.uk/government/publications/generative-artificial-intelligence-in-education/generative-artificial-intelligence-ai-in-education

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

[16] Notice of the General Office of the Ministry of Education on Strengthening AI Education in Primary and Secondary Schools (Document No. 32 [2024] of the General Office of the Ministry of Education's Department of Basic Education; Chinese: 教基厅函〔2024〕32 号) [教育部办公厅关于加强中小学人工智能教育的通知]. Ministry of Education of the People's Republic of China. 2024. http://www.moe.gov.cn/jyb_xwfb/gzdt_gzdt/s5987/202412/t20241202_1165500.html (see also CERNET coverage, "Near-universal AI education in primary and secondary schools targeted before 2030": https://www.edu.cn/xxh/focus/zc/202412/t20241202_2644322.shtml)

[17] *AI Guidance for Schools Toolkit / US States with AI in Education Guidance*. TeachAI (an initiative of Code.org, ETS, ISTE, Khan Academy, and the World Economic Forum, among others). 2024/2025. https://www.teachai.org/toolkit (see also AI for Education's state-level tracker: https://www.aiforeducation.io/ai-resources/state-ai-guidance)

Chapter 7 Assessing AI Literacy: Toward Process-Oriented Evaluation

7.1 Introduction: From "What Was Taught" to "What Was Learned"

Once generative AI entered K-12 classrooms, assessment stopped being the last link in the instructional chain and became the linchpin of the entire AI-education enabling infrastructure. On one hand, AI literacy itself still lacks a widely accepted operational definition or mature measurement instruments — the international community is still arguing over what exactly should be assessed. On the other hand, agentic learning tools can now record, analyze, and even intervene in the learning process in real time, which reopens every downstream question: what to assess it with, when to assess, and who should do the assessing.

This is not an abstract debate; it has real institutional weight behind it. Between 2024 and 2026, the world's major education systems moved almost in lockstep to elevate "AI literacy" from an optional technical interest to a foundational competency — and, increasingly, a legal or policy mandate — for every student. UNESCO published its AI Competency Framework for Students in August 2024, offering the first internationally referenced K-12 framework of its kind [1]. The OECD has designated "Media and Artificial Intelligence Literacy" (MAIL) as an **innovative domain for PISA 2029**, which will test 15-year-olds across more than 90 economies worldwide, with results expected in December 2031 [2][3]. And the EU's Artificial Intelligence Act goes further still: Article 4 writes "AI literacy" directly into law, requiring providers and deployers of AI systems to ensure adequate AI literacy among relevant personnel **as of 2 February 2025** [4]. Assessment is ultimately what will determine whether — and how well — this whole raft of goals gets translated into practice.

Rather than crowning any single instrument as authoritative, this chapter traces the methodological progress, points of contention, and transferable practical frameworks that have emerged globally between 2024 and 2026 across five linked stages: **defining the construct → assessment methods → process-based evidence → academic integrity → equity and governance**.

One position needs to be stated up front: AI literacy assessment is still at an early stage characterized by "many frameworks, few validated scales, still fewer norms, and fewer performance tests still." A systematic review of AI literacy scales identified 22 validation studies covering 16 scales, of which **13 were self-report instruments and only 3 were performance-based tests**. Most of these scales show acceptable structural validity and internal consistency, but **almost none has been tested for cross-cultural validity or measurement error**, and most of the underlying studies neither report interpretability metrics nor make raw data available [5]. Given this, the chapter treats specific reliability/validity coefficients, participating-country or student-sample sizes, and cross-national comparison scores primarily as matters of research framing and mechanism rather than settled fact; any figure this chapter has not independently verified is flagged as [TBD: data/source] rather than risk misleading readers with premature quantitative claims.

> This chapter focuses on assessment as a cross-cutting theme and should be read alongside Chapter 3 (AI curriculum and standards), Chapter 4 (human–AI co-learning), and Chapter 6 (teacher AI literacy). On teachers as assessment agents specifically, see Chapter 6 and our institute's *[TBD: title of dedicated report/chapter on teacher AI literacy]*.

7.2 Defining the Construct: The Prior Question Assessment Cannot Avoid

Every assessment effort begins with defining its construct. The first and most fundamental difficulty in AI literacy assessment is that the construct itself remains open-ended and still shifting.

7.2.1 From Academic Definition to International Frameworks: A Discernible Line of Convergence

The modern operationalization of "AI literacy" is usually traced back to the foundational work of Long and Magerko at the 2020 CHI conference. Drawing on multidisciplinary literature, they proposed that AI literacy comprises **17 core competencies** — from "recognizing AI" (distinguishing artifacts that use AI from those that do not) and "understanding intelligence" (differentiating human, animal, and machine intelligence) through to "ethics" (privacy, employment, misinformation, bias, transparency, and accountability) and "programmability" (understanding that agents can be programmed) — along with a set of design considerations for learners [6]. This definition anchored "critical evaluation" and "responsible use" as the core of AI literacy and became the shared point of origin for nearly every framework that followed.

Since generative AI went mainstream, a wave of institutional frameworks published in 2024–2025 has continued to converge on the Long & Magerko core while expanding it meaningfully in two directions: co-creating with AI, and critically checking generated content. The four most internationally influential frameworks are compared in Table 7-1.

AI Literacy Framework	Issuing Body	Release Date	Core Dimensions (No.)	Progression/Tiers	K-12 Focus	Paired Assessment
What is AI Literacy? (17 competencies)	Long & Magerko (Georgia Tech), CHI 2020	2020	17 competencies + design considerations	No explicit tiers	General (including children)	No (academic review)
AI	UNESCO	2024-08	4 dimensions,	Understand / Apply /	Yes (K-	No (curriculum-

Competency Framework for Students	(Miao, Shiohira, Lao, et al.)		12 competencies	Create	12)	oriented)
AI Literacy: Understand, Evaluate, Use	Digital Promise	2024-06	3 modes of engagement, 6 practices	Three interdependent modes	Yes (PK-12)	No (implementation guide)
Empowering Learners for the Age of AI (AILit)	OECD × European Commission × Code.org	2025-05 (review draft)	4 domains, 22 competencies (KSA)	Knowledge/skills/attitudes throughout	Yes (K-12)	Feeds into PISA 2029

Table 7-1 Comparison of major AI literacy framework dimensions [1][6][7][8]

· **UNESCO's AI Competency Framework for Students (2024-08)** organizes competencies into **four dimensions and 12 competencies** — a human-centred mindset, AI ethics, AI techniques and applications, and AI system design — arranged in three progressive tiers (Understand–Apply–Create), and is positioned as an international reference point for folding AI learning objectives into official curricula [1].

· **Digital Promise's AI Literacy Framework (2024-06)** is built around **three interdependent modes of engagement** — Understand, Evaluate, Use — under which sit six actionable practices: algorithmic thinking, data analysis and inference, data privacy and security, digital expression, ethics and impact, and information and mis/disinformation [7].

· **The AILit framework (OECD × European Commission × Code.org, 2025-05 review draft)** organizes learners' interactions with AI into **four domains — Engage, Create, Manage, and Design — comprising 22 competencies in total**, each elaborated through a knowledge–skills–attitudes structure. The framework is explicitly designed to feed into PISA 2029 MAIL and aligns with the EU's DigComp 2.2 and Article 4 of the EU AI Act [8][4].

Beneath their differing terminology, these frameworks share a **common core** that can be grouped into four clusters:

· **Know & Understand**: grasping AI's basic concepts, data mechanics, capability boundaries, and limitations — being able to tell "is this AI" and "how did AI arrive at this result."

· **Use & Apply**: appropriately selecting and steering AI tools for real tasks, including the newly operationalized sub-skills of prompting and collaboration that generative AI has introduced.

· **Evaluate & Create**: critically judging the reliability, bias, and appropriateness of AI output — the AILit framework explicitly lists the competency of "evaluating whether AI output should be accepted, revised, or rejected" [8] — while also being able to create with AI's help.

· **Ethics & Responsibility**: understanding data privacy, algorithmic bias, academic integrity, and social impact, and developing values consistent with responsible use. Notably, neither

UNESCO's nor AILit's framework treats ethics as a standalone module; instead, it cuts across every dimension as an underlying principle [1][8].

The spread of generative AI has driven two conspicuous expansions of the traditional frameworks: first, **prompting and collaboration** have become independently assessable operational skills (made explicit in AILit's "Create" domain and Digital Promise's "Use" mode); second, **critical verification** — checking AI output for factual accuracy and spotting hallucinations — has moved from a "higher-order skill" down to a foundational competency expected of every student.

7.2.2 Persistent Disagreements Beneath the Convergence

The convergence across frameworks should not obscure three substantive disagreements, each of which directly shapes the boundaries of what gets assessed:

- **Whether to include "designing/building AI."** Both UNESCO and AILit retain "AI system design/designing AI" as a higher-order domain, emphasizing that students should "understand AI design principles from an early age" in order to help shape AI for good [1][8]. Digital Promise, by contrast, leans toward an Understand–Evaluate–Use model for all students, treating system-building as an advanced rather than universal goal [7]. This disagreement determines whether assessment needs to cover hands-on skills such as modeling and training-data selection.
- **"Literacy" versus "competency."** UNESCO and AILit both use "competency/competence" rather than "literacy," emphasizing a transferable, integrated knowledge–skills–attitudes construct; this seemingly small wording choice carries real consequences for how scoring units and evidentiary requirements are defined.
- **Whether — and how — attitudes and mindset can be assessed at all.** Dimensions like a human-centred mindset, curiosity, and a sense of responsibility are near-universally included in these frameworks, yet they are the hardest to measure objectively and remain the epicenter of current debates over scale validity.

7.2.3 From a Knowledge-Based to a Literacy-Based Assessment Orientation

A key methodological judgment underlies this chapter: if AI literacy is narrowed down to mere "facts about AI," assessment will easily slide into something rote and test-friendly — the opposite of what a literacy-based educational goal is supposed to achieve. A literacy-based assessment orientation instead emphasizes **performance in context**: can a student, faced with a real, uncertain task, judge when to use AI, how to steer it, and how to take responsibility for its output? Notably, PISA 2029 MAIL is designed along exactly this line: rather than relying solely on conventional test items, it introduces a "simulated environment involving the internet, social media, and AI tools" so that 15-year-olds generate evidence of their competencies through **simulated real-world interaction** [3].

This orientation is what gives process-oriented assessment and performance tasks — discussed below — their central role.

7.3 The Spectrum of Assessment Methods: Summative and Process-Oriented as Complements

The methods used to assess AI literacy can be arranged along a spectrum defined by their source of evidence, running from standardized tests to embedded learning analytics — each carrying its own validity claims and inherent limitations. Understanding where current practice actually sits on this spectrum matters: as the systematic review cited above shows, existing instruments are **heavily skewed toward self-report (13 to 3)**, meaning that what is currently being measured is largely students' self-perception, not their actual competency [5].

7.3.1 Summative Assessment

- **Self-report scales.** These use Likert-type items to measure students' confidence, attitudes, and perceived competency around AI use, and they are by far the most numerous instrument type. A representative example is the AI Literacy Questionnaire (AILQ) developed by Ng and colleagues, which measures secondary students' AI literacy across **affective, behavioral, cognitive, and ethical** dimensions and showed acceptable reliability and validity in a pilot of 363 students in Hong Kong [9]. Wang and colleagues proposed a 12-item scale spanning **awareness, use, evaluation, and ethics**, validated through factor analysis [10]. Self-report scales are attractive because they scale cheaply and easily, but their fundamental limitation is that **self-reported and actual competency often diverge systematically** — particularly among younger students and novices — and are vulnerable to social-desirability bias.
- **Standardized knowledge and performance tests.** These use multiple-choice or true/false items to gauge conceptual understanding, or tasks with objectively correct answers to gauge actual competency. A methodologically more rigorous example is GLAT (the Generative AI Literacy Assessment Test): a **20-item multiple-choice performance test** developed following classical procedures from psychological and educational measurement, validated in a sample of 355 university students using both classical test theory and item response theory, yielding a stable two-parameter logistic (2PL) model (Cronbach's $\alpha = 0.80$, omega total = 0.81, RMSEA = 0.03, CFI = 0.97). More important still, **GLAT scores significantly predict learners' actual performance on AI-supported tasks, and do so better than self-reported measures such as self-assessed ChatGPT proficiency** [11]. This finding provides direct evidence for the methodological claim that performance tests outperform self-report — though GLAT was built for higher education, and its adaptation and localization for K-12 grade bands remains to be done [TBD: reliability/validity data for a K-12 version].

· **Situational judgment tests (SJTs).** These present dilemmas involving AI use and ask students to choose or rank response strategies, approximating the ethical and decision-making dimensions of literacy — a method that sits between knowledge tests and performance tasks.

7.3.2 Performance-Based and Process-Oriented Assessment

· **Performance tasks.** These assign an open-ended task that genuinely requires invoking AI (for example, "use AI to help complete a research project with source verification"), scored against a rubric that evaluates judgment, collaboration, and verification behavior. The simulated-environment testing used in PISA 2029 MAIL is a large-scale attempt at exactly this [3].

· **Portfolios.** These compile a student's AI-collaboration work and reflections over time, capturing a developmental trajectory rather than a single snapshot. China's *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)* explicitly calls for using digital technology to "build growth records of students' AI literacy" — in essence, a policy-level articulation of the portfolio/growth-record approach [12].

· **Process-based behavioral evidence.** Learning platforms can log prompt histories, revision trails, and multi-turn dialogues with agents, allowing researchers to reconstruct a learner's thinking and decision-making process (discussed further in Section 7.4).

One methodological principle should be held firmly: **no single method can, on its own, support a valid judgment about AI literacy.** Robust assessment design favors triangulation across multiple sources — summative tests to anchor the knowledge base, performance tasks to capture higher-order competencies, and process-based evidence to reveal developmental mechanisms. The gap between GLAT and self-report measures is itself a warning: relying on self-report alone to judge competency risks conclusions that are systematically inflated or deflated [11].

7.4 The Technical Mechanics — and New Possibilities — of Process-Oriented Assessment

The single most profound way that generative AI and agentic tools are reshaping assessment is by making **process-based evidence available** for the first time. This is the central argument of this chapter, and it is also an active research frontier: a bibliometric analysis of "AI-empowered formative assessment" shows the topic has expanded continuously since 2019 and has become an emerging mainstream theme in education research [13].

7.4.1 Learning Processes That Can Now Be Recorded

In a traditional classroom, most of the learning process happens inside a student's head, effectively unobservable — assessment has no choice but to lean on visible outputs. The arrival of AI teaching assistants and AI learning platforms has made a whole category of previously invisible process signals collectible:

· **Prompt-evolution trajectories**: how a student progressively rewrites a question, moving from vague to precise, is itself direct evidence of "use" and "evaluate" competencies.

· **Multi-turn dialogue structure**: whether a student simply accepts AI's first answer or pushes back, questions it, and asks for justification reveals their disposition toward critical verification — mapping directly onto the AILit competency of "evaluating whether AI output should be accepted, revised, or rejected" [8].

· **Revision and adoption behavior**: whether a student keeps, rewrites, or discards an AI suggestion reveals the autonomy of their judgment.

· **Timing and frequency of help-seeking**: at which point in a task chain a student turns to AI reflects metacognition and task-decomposition strategy.

These signals let assessment trace backward from "the artifact" to "the process," offering a level of granularity for formative feedback that was previously unavailable. Recent learning-analytics research has built on this to develop a set of technical approaches for formative assessment, including adaptive feedback, multimodal analysis, predictive modeling, and dashboards [14].

7.4.2 Agents as Providers of Formative Feedback

AI teaching assistants are not just objects of assessment — they can also act as assessors themselves: diagnosing a student's difficulty in real time, offering targeted hints, and feeding process data back to the teacher. Their value lies in compressing the assess–feedback–adjust loop from a weekly cadence down to a matter of minutes, closing in on the ideal of truly formative assessment. Beijing's *Work Plan for Advancing AI Education in Primary and Secondary Schools (2025–2027)* explicitly calls for using AI technology to conduct **learning analytics and homework diagnostics**, building a new intelligent teaching-and-learning model of "teaching guided by learning data, instruction tailored to the individual, and assessment driving instruction" [15] — a concrete example of this formative-assessment mechanism being written into regional policy.

That said, the risks deserve equal billing. A critical review of using learning analytics to support formative assessment systematically maps out this path's **progress, pitfalls, and way forward**, identifying structural gaps including insufficient subject-matter adaptation, low student involvement in feedback design, inadequate attention to ethics, and a limited understanding of generative AI's role in formative assessment [14]. Building on that review, this chapter flags three risks in particular:

· **Process data is not competency.** Frequent interaction does not necessarily mean deep learning, and a tidy-looking trajectory does not necessarily mean genuine understanding; these signals need to be mapped carefully onto the underlying construct, or assessment risks the fallacy of substituting "what's measurable" for "what matters." When learning analytics is used for formative assessment, the validity chain linking its evidence to genuine learning has not yet been fully established [14].

· **The ethical limits of assessment-driven monitoring.** Continuous recording of the learning process can tip over into excessive surveillance of students, raising privacy and psychological-safety concerns that must be governed by the data-governance rules set out in Chapter [TBD: governance chapter number].

· **The risk of feedback dependency.** If instant feedback displaces a student's own opportunities for self-monitoring, it may end up weakening the development of metacognitive skills rather than strengthening it.

7.4.3 Summary Table: Mapping Process-Based Evidence Types to Assessment Uses

Process-Based Evidence Type	Corresponding AI Literacy Dimension	Collection Medium	Main Limitation
Prompt-evolution trajectory	Use / Evaluate	AI learning platform / agent-tutor logs	Hard to distinguish strategic revision from blind trial-and-error
Multi-turn dialogue structure	Evaluate / Critical verification	Conversational agent interaction logs	Automated judgment of dialogue quality remains immature [14]
Revision and adoption behavior	Use / Responsibility	Document revision history / version trail	Motivation behind adoption cannot be directly observed
Reflective writing	Ethics / Metacognition	Portfolio / growth record [12]	Depends on students' expressive ability; limited for younger grades

Table 7-2 Mapping process-based evidence to AI literacy dimensions (collection protocols and automated-scoring reliability/validity await localized development)

7.5 Academic Integrity: A Precondition Challenge for Assessment Validity

The most direct impact of generative AI on assessment systems is that it **undermines the validity of traditional assessment models that treat the final product as the sole evidence.** When an essay, a piece of code, or a report could plausibly have been generated by AI, the inferential chain from "artifact" to "competency" is broken.

7.5.1 The Structural Limits of the Detection Approach

Relying on "AI-generated-content detectors" as a technical safety net has a structural flaw, and there is now hard evidence of it. A Stanford-led study by Liang and colleagues, published in *Patterns*

(Cell Press) in 2023, tested seven mainstream detectors against 91 genuine TOEFL essays written by non-native English speakers: the **false-positive rate reached 61.3%** — more than half of these entirely human-written essays were flagged as AI-generated, and roughly 19% were unanimously misclassified by all seven detectors [16]. The reason this misclassification systematically penalizes non-native writers is that the "AI signatures" these detectors rely on — simple vocabulary, predictable structure, sparse idiom use — happen to overlap with the natural features of second-language writing [16]. In other words, pinning academic integrity on after-the-fact detection is not just unreliable and locked in a losing arms race against evasion techniques — it actively **produces systematic unfairness against specific groups of students**, which makes it untenable as a strategy regardless of how much the technology improves.

Authoritative policy has already pivoted away from this approach for exactly this reason. The U.S. Department of Education's October 2024 *Toolkit for Safe, Ethical, and Equitable AI Integration*, together with accompanying guidance from its Office for Civil Rights (OCR), **recommends against relying on AI detection tools**, citing insufficient reliability and disproportionate disciplinary consequences, and instead advocates a trust-based approach to academic integrity grounded in transparency and process [17].

7.5.2 Redesigning Assessment: From "Preventing" to "Designing For"

A more sustainable path is to **redesign the assessment task itself** so that it can still meaningfully distinguish competency levels even while acknowledging that AI exists:

- **Making process visible.** Extending what gets assessed from the final product to traceable process evidence — drafts, prompt histories, revision trails — reduces the feasibility of having someone else (or something else) do the work.
- **Requiring transparency about AI use.** Making "how AI was used, at which stages, and how it was verified" part of what gets assessed turns AI use from something to hide into a competency that can be evaluated. The UK's Joint Council for Qualifications (JCQ) issued guidance in 2024, *AI Use in Assessments*, requiring that students **label** any part of their work directly generated by AI; using AI without disclosure constitutes malpractice [18].
- **Increasing contextualization and personalization.** Designing tasks tied to local context, personal experience, and in-class events reduces how well generic AI output can substitute for genuine student work.
- **Returning to oral and live performance.** Defenses, live demonstrations, and similar formats verify the genuine understanding that sits behind the process evidence.

Policy boundaries around whether AI use is "permitted / conditionally permitted / prohibited" vary considerably by country and grade band. China's *Guidelines for Generative AI Use by Primary and Secondary Students (2025 Edition)* sets tiered rules by grade level: **elementary students are**

prohibited from independently using open-ended content-generation features, though teachers may use them appropriately in class to support instruction; middle-school students may moderately explore analyzing the logic of generated content; and high-school students are permitted to conduct inquiry-based learning that engages with the underlying technical principles [19]. Table 7-3 compares specific policy texts and their applicable grade bands.

Country/Region	**K-12 Academic Integrity and AI-Use Policy Stance**	**AI-Use Disclosure Required?**	**Stance on Detectors**	**Policy Document/Year**
China	Tiered by grade band: elementary prohibits independent use of open-ended generation, middle school allows moderate exploration, high school allows inquiry-based use	Transparent use encouraged	[TBD: exact document wording]	*Guidelines for Generative AI Use by Primary and Secondary Students (2025 Edition)*, 2025 [19]
United States (federal)	Trust–transparency–process-oriented; recommends against reliance on detectors	Disclosure encouraged; safe space for disclosure promoted	Explicitly recommends against reliance	Department of Education *AI Toolkit* and OCR guidance, 2024-10 [17]
United Kingdom (exam boards)	Undisclosed AI use constitutes malpractice	Students required to label AI-generated portions	Combines investigation with human judgment	JCQ *AI Use in Assessments*, 2024 [18]
European Union	AI literacy codified in law (AI Act, Article 4); teaching-ethics guidance precedes it	—	—	EU AI Act, Article 4, effective 2025-02 [4]

Table 7-3 Comparison of K-12 academic-integrity and AI-use policy stances across major education systems [4][17][18][19]

7.6 Equity and Assessment Bias: Who Gets Assessed, and Who Gets Misjudged

Once AI literacy assessment is deployed at scale, its own fairness becomes a governance issue in its own right. This section highlights three equity risks from a measurement-fairness perspective and proposes practical review thresholds.

- **Unequal access to tools contaminates scores.** If an assessment depends on a student's familiarity with a particular AI tool, and access to that tool itself varies by urban/rural location, school resources, or family socioeconomic status, then the assessment is largely measuring "has this student had exposure" rather than "how literate is this student." This is validity being eroded by the digital divide. (See Chapter [TBD: equity-of-access chapter number] for the structural dimensions of digital access.)
- **Language and cultural bias.** As Section 7.5.1 shows, automated judgments based on generative models or linguistic features can systematically disadvantage non-native speakers, students with dialect backgrounds, and students with special needs — Liang and colleagues' 61.3% false-positive rate for non-native writers is the starkest warning of this [16]. Any automated scoring or situational-judgment item needs differential item functioning (DIF) review. The systematic review of AI literacy scales cited earlier likewise found that **not one existing scale has undergone cross-cultural validity testing**, meaning direct cross-national transplantation carries a real risk of measurement non-equivalence [5].
- **Group-level divergence in self-report bias.** Within self-report scales, the tendency to over- or under-estimate one's own competency differs across gender, grade level, and prior-experience groups; comparing raw self-report scores directly across these groups would widen rather than narrow the resulting gap — which is itself one of the fairness advantages that performance tests such as GLAT hold over self-report scales [11].

The methodological recommendation follows directly: any high-stakes AI literacy assessment used for selection or accountability purposes should first complete a fairness review — covering access preconditions, DIF analysis, cross-cultural invariance testing, and cross-source evidence triangulation — and should not otherwise be used to make high-stakes judgments about students, teachers, or schools. Even the timeline for PISA 2029, which will push AI literacy assessment for 15-year-olds onto a globally unified footing, has drawn public questioning from scholars over whether it is moving too hastily, whether the construct is mature enough, and whether cross-cultural comparability can be trusted [3] — which only underscores the need to put fairness review first.

7.7 An Implementation Framework: Building Assessment Capacity into the Enabling Infrastructure

Bringing the preceding analysis together into actionable components of the enabling infrastructure, a robust K-12 AI literacy assessment system should include the following layers (specific indicators, instruments, and norms for each layer still require localized development; what follows is a structural outline with anchoring real-world reference points):

1. **Construct layer**: Define localized AI literacy dimensions and grade-band progressions rather than importing an external framework wholesale. UNESCO's "Understand–Apply–Create" three-tier progression [1] or AILit's "four domains × knowledge/skills/attitudes" structure [8] can serve as reference points, but should only be adopted after cross-cultural validity testing [5]. The "knowledge–skills–thinking–values" four-dimension structure set out in China's *General AI Education Guidelines (2025 Edition)* can serve as a starting point for a localized construct [12].
2. **Instrument layer**: Combine summative and process-oriented methods, holding firm to multi-source triangulation; prioritize developing performance tests over relying on self-report alone (GLAT's evidence shows performance tests have stronger predictive power) [11]; instruments must report reliability and validity [TBD: localized reliability/validity data].
3. **Data layer**: Collection, storage, and use of process data should be constrained by minimum-necessity and informed-consent principles, consistent with the data rules set out in the governance chapter; the inherent validity and ethical pitfalls of using learning analytics for formative assessment warrant ongoing vigilance [14].
4. **Teacher layer**: Teachers, as the designers and interpreters of assessment, have their own assessment literacy — a bottleneck for system-wide implementation that requires dedicated training. See Chapter 6 and our institute's *[TBD: report related to teacher AI literacy]*.
5. **Governance layer**: Distinguish low-stakes uses (formative, diagnostic) from high-stakes uses (selection, accountability), applying stricter fairness and transparency thresholds to the latter; academic integrity should be built around "designing against cheating" rather than "detecting after the fact" [16][17].

One concrete regional example worth noting is Beijing: starting in fall 2025, all primary and secondary schools citywide have begun offering general AI education, with **no fewer than 8 class hours per academic year**, spanning every grade from elementary through high school; the city has also folded schools' and districts' AI-education implementation into its educational-quality supervision and evaluation system, with regular quality reviews and instructional inspections [15]. This suggests that building assessment capacity is not just a matter of student-facing instruments — it is also an institutional matter of embedding "assessment" into teacher-research and supervisory systems.

7.8 Summary and Research Outlook

This chapter's central judgment is that AI literacy assessment is undergoing a paradigm shift — from product to process, from preventing cheating to evaluating competency — and that generative AI and agentic tools are simultaneously the drivers of this shift and a new source of risk to it. Between 2024 and 2026, international frameworks (UNESCO, Digital Promise, AILit) have converged markedly on the underlying construct [1][7][8]; large-scale assessment (PISA 2029 MAIL) and legal obligation

(Article 4 of the EU AI Act) are pushing AI literacy onto a track where it is both measurable and mandated to be measured [2][3][4]. Yet on the instrument side, the field remains at an early stage — "heavy on self-report, light on performance tests, and lacking in cross-cultural validity" [5] — and the detection-based approach to academic integrity has already been empirically discredited and rejected by authoritative policy [16][17]. The current level of maturity does not support treating any single scale or detection tool as authoritative; the sound approach is to **put the construct first, triangulate across multiple sources, front-load fairness review, and back it all with governance.**

Questions still awaiting answers from the research and policy communities include: how to collect process-based evidence effectively and ethically among younger students; how to adapt performance tests like GLAT for K-12 grade bands and complete localized reliability/validity verification [11]; how to build AI literacy norms that are cross-nationally comparable yet respect local context, validated through cross-cultural invariance testing [5]; and how to rebuild a validity foundation for high-stakes assessment given AI's deep involvement in learning. As relevant cross-national comparison data, scale reliability/validity findings, and adoption figures become available, they will be added to this chapter's tables on an ongoing basis, drawn only from verified real sources.

Chapter References

[1] AI competency framework for students · UNESCO · 2024 (published 2024-08; authors Fengchun Miao, Kelly Shiohira, Natalie Lao; four dimensions, 12 competencies, three-tier Understand/Apply/Create progression) · https://www.unesco.org/en/articles/ai-competency-framework-students

[2] PISA 2029 Media and Artificial Intelligence Literacy · OECD · 2026 (MAIL is an innovative domain for PISA 2029, assessing 15-year-olds, led by ACER, covering 90+ economies, results expected 2031-12) · https://www.oecd.org/en/about/projects/pisa-2029-media-and-artificial-intelligence-literacy.html

[3] Why we should be skeptical of the hasty global push to test 15-year-olds' AI literacy in 2029 · The Conversation · 2025 (describes PISA 2029 MAIL's simulated internet/social-media/AI-tool testing environment, and records scholarly skepticism about the timeline and construct maturity) · https://theconversation.com/why-we-should-be-skeptical-of-the-hasty-global-push-to-test-15-year-olds-ai-literacy-in-2029-263695

[4] Article 4: AI literacy · EU Artificial Intelligence Act · 2025 (Article 4 requires providers and deployers of AI systems to ensure adequate AI literacy, applicable from 2025-02-02) · https://artificialintelligenceact.eu/article/4/

[5] Lintner, T. A systematic review of AI literacy scales · npj Science of Learning · 2024 (22 studies, 16 scales; 13 self-report, 3 performance tests; none tested for cross-cultural validity or measurement error) · https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11303566/

[6] Long, D. & Magerko, B. What is AI Literacy? Competencies and Design Considerations · Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems · 2020 (proposes 17 core AI literacy competencies and design considerations) · https://dl.acm.org/doi/fullHtml/10.1145/3313831.3376727

[7] AI Literacy: A Framework to Understand, Evaluate, and Use Emerging Technology · Digital Promise · 2024 (three modes of engagement — Understand/Evaluate/Use — six AI literacy practices) · https://digitalpromise.org/2024/06/18/ai-literacy-a-framework-to-understand-evaluate-and-use-emerging-technology/

[8] OECD. Empowering Learners for the Age of AI: An AI Literacy Framework for Primary and Secondary Education (Review Draft) · OECD × European Commission × Code.org · 2025-05 (four domains — Engage/Create/Manage/Design — 22 competencies, knowledge–skills–attitudes structure; feeds into PISA 2029) · https://ailiteracyframework.org/wp-content/uploads/2025/05/AILitFramework_ReviewDraft.pdf

[9] Ng, D. T. K., Wu, W., Leung, J. K. L., Chiu, T. K. F. & Chu, S. K. W. Design and validation of the AI literacy questionnaire: The affective, behavioural, cognitive and ethical approach · British Journal of Educational Technology · 2024 (AILQ; four dimensions: affective/behavioral/cognitive/ethical; piloted with 363 secondary students in Hong Kong) · https://bera-journals.onlinelibrary.wiley.com/doi/10.1111/bjet.13411

[10] Wang, B., Rau, P.-L. P. & Yuan, T. Measuring user competence in using artificial intelligence: validity and reliability of artificial intelligence literacy scale · Behaviour & Information Technology · 2023 (12-item scale across four dimensions: awareness/use/evaluation/ethics) · https://ideas.repec.org/a/taf/tbitxx/v42y2023i9p1324-1337.html

[11] Jin, Y., Martinez-Maldonado, R., Gašević, D. & Yan, L. GLAT: The Generative AI Literacy Assessment Test · arXiv:2411.00283 (Monash University) · 2024 (20-item multiple-choice performance test; 355 university students; α=0.80, ω=0.81, RMSEA=0.03, CFI=0.97; predicts task performance better than self-report) · https://arxiv.org/abs/2411.00283

[12] Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [中小学人工智能通识教育指南（2025 年版）] · Ministry of Education Basic Education Teaching Guidance Committee · 2025-05-13 (school-based assessment across "knowledge–skills–thinking–values" four dimensions; performance-based assessment; combining process-oriented and summative assessment; growth records of students' AI literacy; pluralistic assessment mechanisms) · https://www.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[13] AI-empowering Formative Assessment: Research Progress and Future Directions · Proceedings of the 2025 4th International Conference on Artificial Intelligence and Education (ACM) · 2025

(bibliometric analysis shows AI-empowered formative assessment has expanded continuously since 2019 and become an emerging mainstream theme) · https://dl.acm.org/doi/10.1145/3797552.3797589

[14] Banihashem, S. K. et al. A Critical Review of Using Learning Analytics for Formative Assessment: Progress, Pitfalls and Path Forward · Journal of Computer Assisted Learning · 2025 (progress in adaptive feedback/multimodal analysis/predictive modeling/dashboards, alongside gaps in subject-matter adaptation, student involvement, ethics, and understanding of generative AI's role) · https://onlinelibrary.wiley.com/doi/10.1111/jcal.70056

[15] Beijing Work Plan for Advancing AI Education in Primary and Secondary Schools (2025–2027) [北京市推进中小学人工智能教育工作方案（2025—2027 年）] · Beijing Municipal Education Commission · 2025-03 (no fewer than 8 class hours per academic year from fall 2025, covering all grade bands; learning analytics and homework diagnostics; "teaching guided by learning data, instruction tailored to the individual, assessment driving instruction"; folded into educational-quality supervision and evaluation) · https://www.beijing.gov.cn/zhengce/zhengcefagui/202503/t20250310_4029667.html

[16] Liang, W., Yuksekgonul, M., Mao, Y., Wu, E. & Zou, J. GPT detectors are biased against non-native English writers · Patterns (Cell Press) · 2023 (seven detectors show a 61.3% false-positive rate on non-native TOEFL essays; roughly 19% unanimously misclassified) · https://pmc.ncbi.nlm.nih.gov/articles/PMC10382961/

[17] Empowering Education Leaders: A Toolkit for Safe, Ethical, and Equitable AI Integration · U.S. Department of Education (including OCR guidance) · 2024-10 (recommends against relying on AI detection tools due to insufficient reliability and disciplinary consequences; advocates a trust–transparency–process-oriented approach to academic integrity) · https://cdt.org/insights/u-s-department-of-educations-ai-toolkit-and-nondiscrimination-resources-provides-lasting-guidance-for-educators-on-ai-and-civil-rights/

[18] AI Use in Assessments: Your role in protecting the integrity of qualifications · Joint Council for Qualifications (JCQ), UK · 2024 (students must label AI-generated portions; undisclosed use constitutes malpractice) · https://www.jcq.org.uk/knowledge-hub/ai-use-in-assessments-your-role-in-protecting-the-integrity-of-qualifications/

[19] Guidelines for Generative AI Use by Primary and Secondary Students (2025 Edition) [中小学生成式人工智能使用指南（2025 年版）] · Ministry of Education (relevant department) · 2025 (tiered by grade band: elementary students prohibited from independent use of open-ended generation, middle school allows moderate exploration, high school allows inquiry-based use) · https://www.edu.cn/xxh/focus/zc/202505/t20250513_2667992.shtml

Chapter 8 Redefining Enabling Infrastructure: Compute, Data, Platforms, and Client-Side Hardware (AI Glasses / AI Learning Tablets)

8.1 From Supplies to Substrate: A Paradigm Shift in What "Enabling Infrastructure" Means

The parent white paper (SLIBNU, 2022) described K-12 AI-education enabling infrastructure largely in terms of supplies: open educational resources (OER), textbooks, algorithms and datasets, laboratory equipment. That framework took shape before generative AI and the agent paradigm went mainstream, and it carried an implicit assumption — that enabling infrastructure simply meant stocking enough static resources and equipment. By 2026, that assumption no longer captures what is actually happening in classrooms.

One telling piece of evidence is how the policy language itself has shifted. Around 2023, education-informatization documents still centered on "resources, equipment, networks" as their core vocabulary. By April 2026, the *Action Plan for "AI+ Education"* (Document No. 1 [2026] of the Joint Working Group on Science, Technology, and Informatization; Chinese: 教科信〔2026〕1 号) — jointly issued by China's Ministry of Education, the National Development and Reform Commission, the Ministry of Industry and Information Technology, the Ministry of Science and Technology, and the National Data Administration — had unmistakably recentered the discourse on a different, programmable set of elements: "compute, data, models, tools." The plan calls for building a "National Educational AI Computing Service Platform" to "effectively pool AI innovation resources such as compute, data, models, and tools," and for developing "stage-specific educational large language models, with strengthened value alignment, logical reasoning, and safety-ethics capabilities" [1]. That shift signals that the official, national-level definition of enabling infrastructure has moved from a static inventory of supplies to a dynamic, callable substrate of intelligent capability.

The rise of generative models, the always-on presence of agent-based teaching assistants, and the real-time demands of multimodal interaction have together pulled the center of gravity of enabling infrastructure away from "how many resources do we own" and toward "can we deliver a stable, accessible, compliant, programmable layer of intelligent capability." This chapter accordingly redefines enabling infrastructure as four interlocking layers: the **compute layer** (where models run, and at what latency and cost), the **data layer** (where training and operational data come from, how it is governed, and how it flows), the **platform layer** (which packages model capability into orchestrable, auditable teaching services), and the **client-side hardware layer** (the devices learners and teachers physically touch, including AI glasses, AI learning tablets, tablets, and wearable feedback terminals). The health of all four layers together determines whether the literacy, curricula, human–AI collaboration, and assessment practices discussed in earlier chapters can actually be put into practice.

This four-layer substrate is not an abstract theoretical construct of our own; it tracks closely with where the leading international frameworks are heading. UNESCO's *AI Competency Framework for Students* and *AI Competency Framework for Teachers*, both released in September 2024, likewise carve "technology and applications" out from the abstract notion of "AI literacy," explicitly requiring that understanding of AI systems, data, and tools be built into the competency spectrum. The student framework lists 12 competencies across four dimensions — human-centered mindset, AI ethics, AI techniques and applications, and AI system design — while the teacher framework covers 15 competencies across five dimensions [2]. Once "understanding and navigating the AI enabling infrastructure" becomes, in its own right, a core competency certified by an international body, enabling infrastructure stops being a backstage procurement concern and moves squarely onto the stage of educational goals.

> Scope note: This chapter addresses the technology–governance substrate of "enabling infrastructure," which corresponds to the competency–governance dimensions of Chapters 4–7 as an "upper-level capability / lower-level support" pairing. Instructional forms of client-side hardware are covered in depth in this institute's *Global Blue Book on AI Glasses in Education 2026* and *Global White Paper on Educational Robotics Development 2026*.

8.2 The Compute Layer: Device–Edge–Cloud Division of Labor and Education-Specific Constraints

8.2.1 The Logic Behind the Three-Tier Split

- **Cloud**: Hosts training and high-complexity inference for large-parameter foundation models. Highly elastic, but constrained by network conditions, latency, per-call cost, and cross-border data-transfer compliance.
- **Edge**: Deployed in campus server rooms or regional education-cloud nodes, used to cut latency, cache high-frequency requests, and handle local preprocessing and de-identification of sensitive data.
- **Device (on-device)**: Runs lightweight models locally on learning tablets, tablets, and AI glasses, meeting the need for offline operation, low latency, and "data that never leaves the device."

This device–edge–cloud tiering is not an invention of the education sector — it follows an already-established path set by national compute-infrastructure strategy. The October 2023 *Action Plan for

High-Quality Development of Computing Infrastructure* (Document No. 180 [2023] of the Ministry of Industry and Information Technology's General Office of Communications; Chinese: 工信部联通信〔2023〕180 号) — jointly issued by the Ministry of Industry and Information Technology, the Cyberspace Administration of China, the Ministry of Education, the National Health Commission, the People's Bank of China, and the State-owned Assets Supervision and Administration Commission — explicitly calls for "advancing the ubiquitous distribution and coordinated development of cloud-edge-device compute" and "accelerating edge-compute buildout to support low-latency applications," setting a quantitative target of surpassing "300 EFLOPS of compute capacity, with intelligent compute reaching a 35% share" by 2025. The same plan names education as one of six priority sectors, requiring "30 or more benchmark applications" in each [3]. Education was thereby formally written into the "last mile" of national compute-infrastructure planning.

8.2.2 Education-Specific Constraints on the Three-Tier Split

Education settings impose constraints on this three-tier division that differ from consumer and enterprise use cases:

- **Cost sensitivity.** Schools cannot sustainably absorb high-frequency, token-metered usage over the long run. If dozens of students in a single class simultaneously query a commercial cloud API, marginal cost rises linearly with usage frequency — a pattern that is hard to fit into a routine budget.
- **Latency sensitivity.** Classroom interaction, spoken-language practice, and real-time annotation all have hard response-time requirements — once end-to-end latency exceeds roughly one second, the sense of natural "conversation" breaks down. This is precisely what motivates on-device processing: when Qualcomm demonstrated the Snapdragon AR1+ Gen 1 platform running Llama 3.2 1B on-device at AWE 2025, the disclosed "time to first token" was about 1.2 seconds [4] — a real, publicly documented reference point for gauging what counts as an acceptable classroom-latency order of magnitude.
- **Compliance sensitivity.** Cross-border transfer and export of minors' data face tighter restrictions (see Section 8.3), which further pushes sensitive-data processing toward the device and edge on legal grounds.

Together, these three constraints push compute "downward" toward the device and edge — converging, directionally, with the national top-level design for "ubiquitous cloud-edge-device coordination."

8.2.3 The Technical Drivers Behind On-Device Processing

The improved capability of small on-device models stems from the combined progress of model compression, quantization, distillation, and dedicated inference chips (NPUs), which together let mid-sized models run at acceptable latency on mobile-class compute. Qualcomm's Snapdragon AR1+ Gen 1, launched in June 2025, is illustrative: powered by a third-generation Hexagon NPU, the platform can run a roughly 1-billion-parameter (1B) Llama 3.2 model with 128K context length directly on smart glasses, with no phone or internet connection required [4]. Academic work echoes this trend — several 2025 systems papers optimizing LLM inference for mobile NPUs report substantial speedups in both the prefill and decoding stages, evidence that running 1B–10B-parameter models on mobile-class compute is moving from demo to engineering-ready practice [5].

For education specifically, the value of on-device processing is not just bandwidth savings — it is moving the "privacy wall" forward to the device itself. Sensitive student voice, responses, facial data, and work products can be structured and reasoned over locally, with only the minimum necessary information sent upstream. This aligns with the "local, closed-loop data" orientation that recurs throughout this institute's client-side hardware research (see the relevant chapters of the institute's *Global Blue Book on AI Glasses in Education 2026*).

8.2.4 Comparing the Three Compute Tiers in Education

Deployment Tier	Latency Profile	Per-Call Cost	Privacy/Compliance Profile	Typical Education Use	Real-World Reference
Cloud	Volatile, shaped by network conditions and queuing	Billed per token/call; scales linearly with frequency	Data may cross borders; requires compliance review	Large-model grading, content generation, teaching-research analytics, training of education-vertical LLMs	National Educational AI Computing Service Platform, education-vertical LLMs [1]
Edge	Low latency; closed loop within campus/region	One-time or amortized cost; low marginal cost	Closed loop within campus/region; de-identified preprocessing	Local caching, de-identified preprocessing, regional sharing	National "cloud-edge-device coordination" edge-compute buildout [3]
Device	Lowest latency; works offline	Near-zero marginal cost	Data can stay entirely on-device	Offline tutoring assistance, real-time interaction, wearable feedback	AR1+ on-device 1B model, on-device large models in AI learning tablets [4][10]

Note: This table characterizes the relative features and real-world reference points of each tier. Actual latency and cost figures depend heavily on network conditions, device model, model scale, and concurrency, and this Blue Book does not fabricate a uniform range for them; the "real-world reference" citations are all independently verifiable public facts.

8.3 The Data Layer: From "Dataset as Supply" to "Data Governed as a Lifecycle"

The parent white paper treated data as a countable supply item — "algorithms and datasets." The reality in 2026 is that data flows continuously across training, fine-tuning, retrieval-augmented generation (RAG), operational logs, and assessment records. The critical question for enabling infrastructure is no longer "do we have a dataset," but "is data governed compliantly and auditably across its entire lifecycle."

8.3.1 Minors' Data: A Legally Mandated "High-Protection" Category

Under Chinese law, minors' data receives explicit heightened protection. Article 28 of the *Personal Information Protection Law of the People's Republic of China* (PIPL) classifies the personal information of minors under age 14 directly as **sensitive personal information**. Article 29 requires that processing sensitive personal information obtain the individual's **separate consent**. Article 31 further requires that processing personal information of minors under 14 obtain the **consent of a parent or other guardian**, and that processors **establish dedicated rules for handling such information** [6]. In practice, this means any AI teaching tool aimed at young students already falls within a fourfold compliance threshold — sensitive personal information, separate consent, guardian consent, and dedicated processing rules — before it ever collects a single voice sample, image, response, or behavioral trace.

International regulation is tightening in the same direction. The U.S. Federal Trade Commission (FTC) finalized amendments to the *Children's Online Privacy Protection Rule* (COPPA Rule) in January 2025, requiring operators to obtain separate, verifiable parental consent before using information collected from children for "non-core" purposes such as targeted advertising, and explicitly prohibiting indefinite retention of children's data — retention is now limited to the reasonable period needed to fulfill the purpose for which the data was collected [7]. The EU's *Artificial Intelligence Act* (EU AI Act) takes a product-classification approach: it designates education as a **high-risk** domain (covering uses such as admissions decisions, learner assessment, and exam proctoring), which triggers strict compliance obligations. More consequentially, the Act classifies **emotion recognition** — inferring student emotion from facial expression, voice, or other biometric signals within educational institutions — as an **unacceptable risk** and **bans it outright**.

That ban has applied since February 2025, while compliance obligations for high-risk education systems apply from December 2027 [8]. Together, these provisions draw a clear legal red line around contested applications such as "classroom emotion AI" and "attention monitoring."

8.3.2 Four Threads of Data Governance

- **Provenance and authorization**: The copyright and usage-authorization boundaries of textbooks, question banks, and student work used for training and retrieval should be governed by a clear chain of authorization with defined use limits.
- **Protection of minors**: Data minimization, purpose limitation, retention limits, and a working right-to-be-forgotten mechanism must align with PIPL Articles 28/29/31 and equivalent protections for minors in other jurisdictions [6][7].
- **Data flows and cross-border transfer**: Cross-border training and cloud-based calls trigger data-export review, which is a direct compliance driver pushing sensitive data "downward" to the edge and device (echoing Section 8.2).
- **Explainability and audit trails**: For assessment and academic-integrity purposes, operational data must be traceable (echoing the assessment and governance discussion in Chapters 6–7).

8.3.3 Two New Categories of Governance Object Introduced by the Generative Paradigm

The generative paradigm introduces two categories of data-governance object that had no counterpart in the parent framework:

- **Prompts and conversation logs**: These may contain student privacy, family information, and other sensitive content. Chinese governance already provides an institutional interface here: generative AI services with public-opinion attributes or social-mobilization capacity that are offered to the public must be **filed/registered** as required by law. According to filing information disclosed by the Cyberspace Administration of China, several hundred generative AI services had completed filing by the end of 2025, with several hundred more applications or features completing registration [9]; education-oriented services must operate within this "filing–security assessment" framework.
- **Model-output audit trails**: Needed for fact-checking, bias auditing, and appeals review. The risk of "hallucination" and inappropriate content from generative models makes output-side record-keeping a necessary piece of infrastructure for academic integrity and content safety.

These two categories of "new data" are governance blind spots that any 2026 rebuild must close, and they are the data-side precondition for making the platform layer's "auditability" requirement (Section 8.4) concrete.

8.4 The Platform Layer: Orchestrating Model Capability into Auditable Teaching Services

Compute and data are of little use to teachers unless packaged by a platform. The platform layer's job is to turn underlying model capability into services that are **classroom-facing, orchestrable, governable, and auditable**.

8.4.1 National Platforms: From "Resource Repository" to "Intelligent Substrate"

China's National Smart Education platform offers a good case study for observing this evolution quantitatively. As of December 2025, the National Smart Education Public Service Platform had surpassed 178 million total users, more than 72.6 billion cumulative visits, roughly 52 million average daily visits, and coverage spanning more than 200 countries and regions [11]. Its accumulated resource base includes more than 130,000 quality K-12 resources, more than 12,500 flagship vocational-education courses, and 145,000 quality higher-education courses, and it has organized more than 500,000 teacher research groups and trained more than 90 million teachers cumulatively [11]. More significant still is the shift in the platform's underlying form: the "National Smart Education Platform 2.0 Intelligent Edition," launched in March 2025, introduced intelligent interaction, knowledge graphs, and multimodal data-analysis capability — marking a transition of the national platform from a "resource-retrieval repository" to an "intelligent capability substrate" [12]. The *Action Plan for "AI+ Education"* takes this further by deploying a "National Educational Big Data Center" (building a cross-department, cross-region, cross-platform data network) alongside the "National Educational AI Computing Service Platform," jointly packaging compute, data, models, and tools into a public service callable by schools at every level [1].

8.4.2 Key Capabilities of the Platform Layer

1. **Model orchestration and routing**: Selecting the appropriate model across the device–edge–cloud continuum for a given task, balancing cost, latency, and compliance (building on the three-tier division in Section 8.2).
2. **Agent and tool orchestration**: Combining teaching-assistant agents, retrieval, grading, and generation capabilities into instructional workflows (building on the human–AI collaboration discussion in Chapter 5).
3. **Safety and compliance guardrails**: Content-safety filtering, age adaptation, protection against unauthorized calls, and audit logging — directly corresponding to the requirement in the *Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition)* [中小学生成式人工智能使用指南（2025 年版）] that use be "premised on protecting personal privacy

and data security" and graded by school stage (elementary students are barred from independently using open-ended content-generation features; teachers may use such tools in class as appropriate instructional support) [13].

4. **Identity and permissions**: Distinguishing the capability boundaries and data-visibility scope of teachers, students, and administrators.
5. **Accessibility adaptation**: Multilingual support, accessibility features, and low-bandwidth degradation (building on the equity-of-access discussion in Chapter 9).

8.4.3 From "Black-Box Calls" to "Auditable Services"

The governance value of the platform layer lies in converting "black-box calls" into "auditable services": every generation and every judgment should be traceable back to a model version, a data source, and a guardrail policy, so as to support regulatory oversight, appeals, and continuous improvement in teaching and research. This requirement is increasingly being anchored institutionally. The *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)* [中小学人工智能通识教育指南（2025 年版）] explicitly calls for "developing data-security management standards for AI education data that specify security requirements for data collection, storage, transmission, and use," and for "establishing privacy-protection mechanisms and regulating admission of AI teaching tools and products" [14]. The *Action Plan for "AI+ Education"* separately requires strengthening the "AI+ education" security-protection system, implementing filing mechanisms for algorithms and large models, and exploring algorithm-security assessment [1]. Today, mainstream education AI platforms on the market vary considerably in capability coverage and compliance maturity, and a unified interoperability and evaluation standard is still lacking [TBD: authoritative evaluation/standard source] — a gap that is precisely where group standards and third-party evaluation have a role to play.

8.5 The Client-Side Hardware Layer: AI Glasses and AI Learning Tablets as Instructional Vehicles

Client-side hardware is the one layer of enabling infrastructure that learners and teachers touch **directly**, and its form factor directly shapes the interaction paradigm and the boundaries of accessibility. Two categories of device are especially notable in 2026:

8.5.1 AI Learning Tablets: The Most Widely Adopted On-Device Agent Terminal

As the most widely adopted client-side device, AI learning tablets are evolving from "content players" into "on-device agent terminals." Market data confirms both the scale and the growth rate of this category. According to RUNTO (洛图科技), China's all-channel sales of learning tablets reached

5.923 million units in 2024, up 25.5% year over year, with sales revenue of RMB 19.06 billion, up 37.6% year over year [15]. According to IDC, China's learning-tablet shipments reached 1.54 million units in Q2 2025, up 44.6% year over year [16]. Market concentration is high: the top five vendors together held about 82.3% share in Q2 2025, with Zuoyebang (作业帮) at approximately 28.5%, iFLYTEK (科大讯飞) at approximately 20.9%, Xueersi (学而思) at approximately 19.3%, and Step by Step (步步高) at approximately 10.1% [16]. All the leading brands market "large model plus hardware" as their core value proposition.

The key technical inflection point is the arrival of **on-device large models**. iFLYTEK's "Spark on-device large model," built for smart hardware, is designed for privacy-sensitive, offline scenarios: data processing happens entirely on the device, with no need to upload to the cloud, reducing leakage risk [10]. Its learning-tablet product line has also integrated iFLYTEK Spark and DeepSeek models to deliver human-like essay grading, interactive math tutoring, and English-speaking practice [10]. This has turned "offline Q&A support, personalized practice, and privacy-first learning analytics" from marketing language into deliverable capability. Key adaptation issues that remain include eye-health and screen-time management, parental controls, and the continuing capability gap between on-device and cloud models [TBD: authoritative benchmark comparing on-device and cloud capability].

8.5.2 AI Glasses and Wearable Feedback Terminals: An Interaction Frontier That Is Also an Ethical Frontier

AI glasses are defined by first-person (POV) multimodal sensing and near-eye interaction, introducing a new interaction and data-capture paradigm for the classroom — while also amplifying concerns about privacy, attention, and ethics.

The market is in a steep growth phase. According to IDC, global smart-glasses shipments reached about 4.065 million units in the first half of 2025, up about 64.2% year over year, with Meta holding more than 70% share; IDC projects global smart-glasses shipments will exceed 40 million units by 2029 [17]. China's market is expanding just as fast, with smart-glasses shipments rising from about 680,000 units in 2023 to about 1.33 million units in 2024, an increase of roughly 95.6% [18]. The key hardware enabler is the on-device AI chip: Qualcomm's Snapdragon AR1+ Gen 1 makes running a 1B-parameter language model entirely on-device — with no phone required — a shipping reality, while shrinking die area by about 26% and cutting power draw on key workloads by about 7% relative to the prior generation [4], clearing some of the obstacles to long-wear comfort and lightweight temple design.

But the ethical and regulatory constraints on classroom use are unusually strict. AI glasses inherently support covert photography, audio recording, and network connectivity, and have already sparked real-world controversy in contexts such as exam cheating [19]. Meanwhile, the EU's ban on emotion recognition in education (Section 8.3) and China's strong statutory protection of minors' sensitive personal information together impose a hard constraint on any notion of "routinely wearing recording devices in the classroom." A systematic study of the instructional forms, timing, and feedback mechanisms of such devices appears in this institute's *Global Blue Book on AI Glasses in Education 2026* and related patent and group-standard work. From the standpoint of enabling infrastructure, this class of device pushes the "privacy wall" and real-time-performance requirements to their limit: multimodal data should be structured and de-identified **on-device or at the edge**, with only the minimum necessary information sent upstream, in order to satisfy both minors' data-protection requirements and classroom-latency constraints. The structured representation of related multimodal learning-behavior data and the associated interaction-feedback protocol should reference the group standard currently being advanced by this institute [TBD: standard number/project-filing source].

8.5.3 An Analytical Framework: Principles for Selecting Client-Side Hardware in Education

Drawing together the legal, market, and technical facts above yields five principles for school procurement and governance:

- **Accessibility first**: Device cost and maintainability must work across urban–rural and other group differences (building on Chapter 9). Learning tablets already have a solid inclusive-access base; AI glasses remain at an early stage in both price and form factor.
- **Privacy on-device by default**: Sensitive data should be processed on-device by default, with upstream transmission held to the minimum necessary — this is both a compliance requirement (PIPL Articles 28/29/31 [6]) and, increasingly, technically feasible (on-device large models [10], on-device 1B models [4]).
- **Manageable and controllable**: Manageability, auditability, and the ability to shut a device down, accessible to both schools and parents, corresponding to the stage-graded control requirements in the generative-AI usage guidelines [13].
- **Health and ethics**: Screen/usage time, visual health, attention, and social effects should all be part of procurement evaluation; emotion-recognition features require particularly careful attention to their legal boundaries [8].
- **Interoperability**: Avoiding vendor lock-in by following open data and interface standards [TBD: source for education client-side hardware interoperability standards].

8.6 Summary: Enabling Infrastructure as the Physical Precondition for "Capability–Governance" to Take Hold

This chapter has reconstructed enabling infrastructure from the parent white paper's inventory of "resources–textbooks–algorithms–laboratories" into a four-layer technology–governance substrate: **compute–data–platform–client-side hardware**. Its central claim is that, under the generative and agent paradigm, the competitiveness of enabling infrastructure no longer depends on how many static resources an institution owns, but on whether the four layers can jointly deliver **low-latency, low-cost, privacy-compliant, auditable** intelligent capability — while adapting downward to equity of access and supporting upward the literacy and assessment goals discussed earlier in this book.

A clear causal chain runs through the four layers. Cloud-edge-device coordination at the compute layer (per the national compute action plan [3]) creates the physical conditions for sensitive data to move downward; strong protection rules at the data layer (PIPL, COPPA, the EU AI Act [6][7][8]) in turn reinforce the push toward compute decentralization and on-device processing; the platform layer packages both into auditable services (National Smart Education Platform 2.0, the filing mechanism [11][12][9]); and the client-side hardware layer delivers the entire substrate into the hands of teachers and students (the scaling of learning tablets and AI glasses [15][16][17]). A shortfall at any one of the four layers — insufficient on-device compute, gaps in data compliance, an unauditable platform, or inaccessible hardware — becomes a real-world bottleneck for the capability–governance goals discussed in earlier chapters.

Countries differ systematically in how they invest across these four layers and in their policy orientation. China's approach centers on public supply — "national platform + education-vertical LLMs + downward-shifted compute" [1][3][11]; the EU is characterized by compliance-first, risk-tiered regulation, including its emotion-recognition ban [8]; and UNESCO, through its competency frameworks, has established "navigating the enabling infrastructure" as a universal educational goal [2]. Quantitative comparison and evidence-based visualization of these differences will be developed further in later chapters and in this institute's ongoing series of studies.

Chapter References

[1] *Action Plan for "AI+ Education"* (Document No. 1 [2026]; Chinese: 教科信〔2026〕1 号)

[《"人工智能+教育"行动计划》] · China's Ministry of Education, National Development and

Reform Commission, Ministry of Industry and Information Technology, Ministry of Science and Technology, National Data
Administration · 2026 · https://www.chsi.com.cn/jyzx/202604/20260410/2293459694.html (also see

the Ministry of Education portal:
http://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html)

[2] UNESCO AI Competency Frameworks for Students and for Teachers · UNESCO · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-students ; https://www.unesco.org/en/articles/ai-competency-framework-teachers

[3] *Action Plan for High-Quality Development of Computing Infrastructure* (Document No. 180 [2023]; Chinese: 工信部联通信〔2023〕180 号) [《算力基础设施高质量发展行动计划》] · Ministry of Industry and Information Technology, Cyberspace Administration of China, Ministry of Education, National Health Commission, People's Bank of China, State-owned Assets Supervision and Administration Commission · 2023 · https://www.edu.cn/xxh/focus/zc/202310/t20231010_2512506.shtml

[4] Qualcomm announces smaller Snapdragon AR1+ Gen 1 chip for smart glasses (AWE 2025; on-device Llama 3.2 1B, about 1.2-second time to first token, Hexagon NPU, 128K context, about 26% smaller die / about 7% lower power draw) · Qualcomm / 9to5Google · 2025 · https://9to5google.com/2025/06/10/snapdragon-ar1-gen-1/

[5] Fast On-device LLM Inference with NPUs (systems research on accelerating LLM inference on mobile NPUs) · ASPLOS 2025 / arXiv · 2024–2025 · https://arxiv.org/pdf/2407.05858

[6] *Personal Information Protection Law of the People's Republic of China*, Articles 28, 29, and 31 [《中华人民共和国个人信息保护法》第二十八条、第二十九条、第三十一条] (personal information of minors under 14 classified as sensitive personal information; processing sensitive personal information requires separate consent; requires guardian consent and dedicated processing rules) · Standing Committee of the National People's Congress · 2021 · https://www.spp.gov.cn/spp/fl/202108/t20210820_527244.shtml

[7] FTC finalizes amendments to the Children's Online Privacy Protection Rule (COPPA Rule) (January 2025; separate verifiable parental consent required for non-core uses; indefinite retention prohibited) · U.S. Federal Trade Commission (FTC) · 2025 · https://www.secrss.com/articles/75025

[8] EU AI Act: education classified as high-risk; emotion recognition in educational institutions banned as unacceptable risk; took effect August 2024, emotion-recognition ban applies from February 2025, high-risk education-system obligations apply from December 2027 · European Union / EU Artificial Intelligence Act · 2024 · https://artificialintelligenceact.eu/high-level-summary/ ; https://artificialintelligenceact.eu/annex/3/

[9] Cyberspace Administration of China announcement on filed generative AI services (several hundred services filed and several hundred applications/features registered as of end of

2025) · Cyberspace Administration of China · 2025–2026 · https://www.cac.gov.cn/2026-01/09/c_1769688009588554.htm

[10] Spark on-device large model / iFLYTEK learning-tablet large-model capability [星火本地大模型 / 科大讯飞学习机大模型能力] (on-device local processing, no cloud upload required; iFLYTEK Spark and DeepSeek integrated into learning tablets) · iFLYTEK · 2025 · https://www.xfyun.cn/solutions/local_model ; https://www.qbitai.com/2025/06/300819.html

[11] National Smart Education Public Service Platform surpasses 178 million total users, 72.6 billion cumulative visits, about 52 million average daily visits, coverage of more than 200 countries and regions, plus resource base and teacher-training scale · Xinhua / Ministry of Education · 2025 · http://www.moe.gov.cn/fbh/live/2025/77791/mtbd/202512/t20251231_1425330.html

[12] National Smart Education Platform 2.0 Intelligent Edition launched (intelligent interaction, knowledge graphs, multimodal data analysis) · Guangming Online · 2025 · https://news.gmw.cn/2025-03/29/content_37935751.htm

[13] *Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition)* [《中小学生成式人工智能使用指南（2025 年版）》] (premised on protecting personal privacy and data security; elementary students barred from independently using open-ended content-generation features, teachers may use such tools as appropriate in-class instructional support) · Ministry of Education Basic Education Teaching Guidance Committee · 2025 · http://www.duozhi.com/industry/insight/2025051217271.shtml

[14] *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)* [《中小学人工智能通识教育指南（2025 年版）》] (developing data-security management standards for AI education data; establishing privacy-protection mechanisms and regulating admission of AI teaching tools and products) · Ministry of Education Basic Education Teaching Guidance Committee · 2025 · https://www.eol.cn/zhengce/wenjian/202505/t20250512_2667827.shtml

[15] RUNTO: China's all-channel learning-tablet sales reached 5.923 million units in 2024 (up 25.5% year over year), with sales revenue of RMB 19.06 billion (up 37.6% year over year) · RUNTO (洛图科技) / 199IT · 2024 · https://www.199it.com/archives/1738340.html

[16] IDC: China's learning-tablet shipments reached 1.54 million units in Q2 2025 (up 44.6% year over year), with top-five vendor shares (Zuoyebang 28.5%, iFLYTEK 20.9%, Xueersi 19.3%, Step

by Step 10.1%) · IDC / Sina Finance · 2025 · https://finance.sina.com.cn/tech/digi/2025-09-04/doc-infphyzn3583924.shtml

[17] Global smart-glasses shipments reached about 4.065 million units in H1 2025 (up about 64.2% year over year), with Meta holding more than 70% share; IDC projects more than 40 million units globally by 2029 · IDC / Wall Street CN · 2025 · https://wallstreetcn.com/articles/3753479

[18] China's smart-glasses shipments rose from about 680,000 units in 2023 to about 1.33 million units in 2024 (up about 95.6% year over year) · Qianzhan Industry Research Institute · 2025 · https://www.qianzhan.com/analyst/detail/220/250626-249e8fd6.html

[19] Controversy over covert photography and cheating with AI glasses in education/exam settings · 36Kr · 2025 · https://www.36kr.com/p/3705588080505221

Chapter 9 Teacher AI Literacy and Professional Development

Fig. 3 The rich-poor gap in AI training and guidance

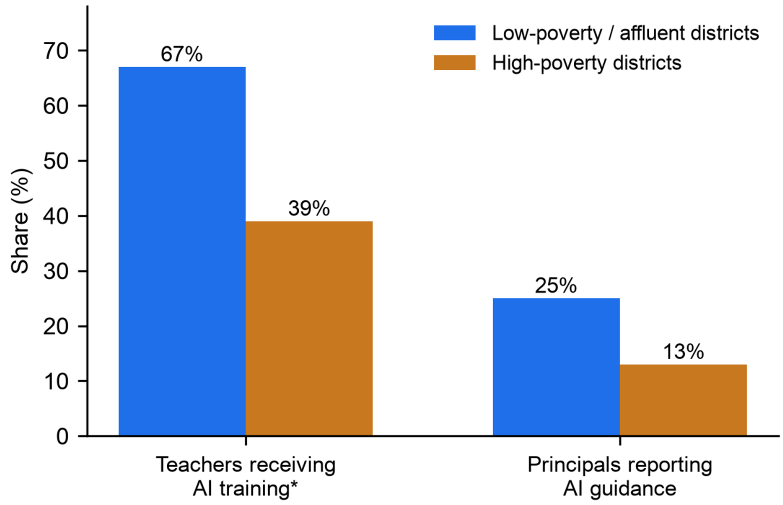

Source: RAND Corporation / TeachAI (2024-2025). * Samples and definitions differ across surveys; figures are not directly additive.

9.1 Why Teacher Literacy Is the Pivotal Variable in Enabling Infrastructure

Within this Blue Book's capability–governance framework, teachers are not passive users of an AI education enabling infrastructure. They are the **hinge node** that connects policy design, technical capability, and student learning outcomes. Policy documents can set curriculum standards; platforms can supply tools and resources. But whether any of that actually converts into classroom learning gains depends on the teacher at the front of the room and the professional judgment she brings. Every thread running through the earlier chapters of this Blue Book — the literacy framework in Chapter 4, curriculum implementation in Chapter 5, human–AI collaboration in Chapter 7, learning assessment in Chapter 8 — ultimately converges on the same constraint: does the teacher have the literacy to make it work?

Since generative AI entered public view in late 2022, the nature of that constraint has shifted structurally. The original master white paper (SLIBNU, 2022) predates this shift, and its discussion of teachers centered on the use of open educational resources (OER) and participation in online teacher communities — at the time, "teacher AI literacy" was still largely a matter of tool operation and resource retrieval. By 2024–2026, three converging technology trends — agentic AI, multimodality, and on-device computing (discussed at length in this institute's *Global Blue Book on

AI Glasses in Education 2026* and *Global White Paper on Educational Robotics 2026*) have pushed teachers into a new role that is **collaborator and governor in equal measure**: teachers now need to divide labor with agentic teaching assistants, catch and correct errors in multimodal generated content, and shoulder frontline responsibility for data and ethics in classrooms where on-device hardware is becoming the norm.

International authorities have already confirmed this shift in role. When UNESCO released its *AI Competency Framework for Teachers* alongside its *AI Competency Framework for Students* in September 2024, it stated plainly that most countries have yet to define the AI competencies expected of teachers or to build national training programs to deliver them — "leaving many educators without proper guidance" [1]. That diagnosis comes from a supranational body, and it matters: the institutional absence of teacher AI literacy is a global shortfall, not a localized one.

This chapter's argument follows from that starting point: the lag in teacher AI literacy has become the most widespread — and most easily underestimated — weak point in the global K-12 AI education enabling infrastructure. This is not a qualitative hunch; it rests on primary survey data. In the United States, RAND Corporation's nationally representative panel of teachers and school districts found that as of spring 2024, roughly seven in ten K-12 teachers had received no training whatsoever on using AI in the classroom [2][3]. Even after training expanded rapidly through fall 2024, only 43% of teachers reported having received AI-related training at least once [2][3]. In other words, nearly two years after AI entered the public consciousness, most frontline teachers remain effectively untrained — which means every other investment the enabling infrastructure makes, on the policy, platform, and resource side, risks losing most of its value in the "last mile."

9.2 The Competency Structure of Teacher AI Literacy

Before asking what literacy teachers need, we need a reference structure to discuss it against. This section first introduces the UNESCO framework, now the international public benchmark, and then presents this institute's four-dimensional model — built to align with that framework while following the "literacy → curriculum → human–AI collaboration → assessment → governance" thread that runs through this Blue Book.

9.2.1 The International Public Benchmark: UNESCO's AI Competency Framework for Teachers

UNESCO's *AI Competency Framework for Teachers* (released September 2024; formally published August 2024) is structured as a two-dimensional matrix of **five dimensions × three progression levels**, yielding 15 competency units in total [1][4]. The five dimensions are: a human-centred mindset, ethics of AI, AI foundations and applications, AI pedagogy, and AI for professional learning [1][4]. The three progression levels are Acquire, Deepen, and Create [1][4]. The framework is anchored in three foundational principles — protecting teachers' rights, strengthening human agency, and promoting sustainability [1][4] — and is positioned as a **global reference tool to guide countries in developing their own teacher competency frameworks, training programs, and assessment parameters** [1][4].

It's worth noting how this framework relates to UNESCO's earlier *ICT Competency Framework for Teachers* (ICT CFT, comprising 18 competencies): the lineage is direct, but the increment is substantive [5]. Where the ICT CFT focused on teachers' mastery and integration of information-technology tools, the AI Competency Framework is the first to elevate "ethics" and "human-centredness" to independent dimensions on par with technology and pedagogy — a reflection of how much more the generative-AI era demands of teachers' **judgment and value-based gatekeeping**.

9.2.2 This Institute's Four-Dimensional Model

While keeping it mappable onto UNESCO's five dimensions, this institute decomposes teacher AI literacy into four dimensions that build on one another in sequence. What follows is an analytical framework; the specific assessment norms and proficiency thresholds for each dimension remain empirical questions to be filled in with real data.

(1) Foundational understanding — corresponding to UNESCO's "AI foundations and applications." This is a teacher's grasp of how AI works at a basic level, where its capabilities end, and how it typically fails. It does not require teachers to become algorithm experts, but it does require them to explain why a generative model can produce output that looks plausible but is in fact wrong (hallucination), and to judge whether a given task is even suited to AI in the first place. Without this dimension, teachers tend to swing between two extremes: blind trust in model output, or wholesale rejection driven by fear.

(2) Pedagogical integration — corresponding to UNESCO's "AI pedagogy." This is the ability to embed AI tools into subject-specific teaching while preserving pedagogical soundness. What matters is not how many AI tools get used, but **whether their introduction serves a clear learning objective.** This dimension connects directly to the classroom models of human–AI collaboration discussed in Chapter 7: teachers must decide which parts of a lesson an agentic teaching assistant can handle, and which parts must stay in human hands to preserve the human, formative dimension of teaching.

(3) Prompt engineering and content stewardship — a core competency that generative AI has newly introduced, sitting within the "Deepen/Create" levels of UNESCO's "AI foundations and applications" dimension. It covers designing effective prompts aimed at instructional goals, fact-checking and vetting generated content for pedagogical fitness, and catching bias or inappropriate material across multimodal outputs — text, image, audio, and video alike. This dimension is a pure addition relative to the original master white paper, which did not address it at all, and it is also where current teacher-training demand is most concentrated.

(4) Ethics and governance — corresponding to the two dimensions UNESCO elevated in parallel: "ethics of AI" and "human-centred mindset." This covers a teacher's frontline ability to uphold academic integrity, protect data, ensure algorithmic fairness, and safeguard student rights. Teachers are the last link in the enabling infrastructure's governance chain: how student data gets collected, where the line sits between AI-assisted work and academic dishonesty, whether disadvantaged groups are being systematically overlooked by algorithmic systems — all of these governance principles are only as real as the daily judgment calls teachers make on the ground. This dimension couples directly with the governance framework in Chapter 10 and the equity-of-access discussion in Chapter 6.

> **Framework summary**: These four dimensions are not a parallel checklist — they build on each other in sequence. Foundational understanding is the prerequisite, pedagogical integration is the main thread, prompt engineering is the operational core of the new era, and ethics and governance is the constraint running through all of it. This layered structure is compatible with UNESCO's "Acquire–Deepen–Create" progression logic: training that targets a single dimension or a single level, in isolation, rarely succeeds on its own.

9.3 Comparing Teacher Competency Standards and Training Policy Across China, the United States, and the EU

How institutionalized teacher AI literacy is serves as a useful proxy for how mature a country's enabling infrastructure has become. Following this Blue Book's standard cross-national comparative format, this section draws on primary policy documents and official data from 2024–2026 to map how China, the United States, and the European Union each approach teacher AI competency standards and training investment.

9.3.1 Comparative Table Across the Three Jurisdictions

Dimension	**China**	**United States (federal level)**	**European Union**
Teacher AI competency standard/framework	Under the "AI+ Education" Action Plan, China will establish a national teacher AI literacy standard and build a tiered training and assessment system matched to different teaching roles [6][7]	No unified federal teacher AI competency standard; relies primarily on the UNESCO framework and standards set independently by states and districts (state-level guidance has been issued by [TBD: number of states], [TBD: source])	Under the *European Framework for the Digital Competence of Educators* (DigCompEdu), AI-related competencies are incorporated; also supported by the *Ethical Guidelines on the Use of Artificial Intelligence and Data in

			Teaching* (first edition 2022, revised alongside the EU AI Act) [8][9]
Included in teacher licensure/registration requirements	Explicitly incorporates "AI knowledge into teacher qualification exams and certification" [6][7]	No (no federal requirement; [TBD: whether any state-level requirement exists, and source])	No; competency framework and ethical guidelines serve as advisory guidance rather than mandatory licensure requirements [8]
National/federal training program	Through the National Smart Education Platform and related channels, more than 2.97 million teachers had received AI literacy training as of end-2024 [10][11]	Executive Order 14277 requires the Secretary of Education to prioritize teacher AI training within discretionary grant programs within 120 days; NSF has concurrently created teacher training opportunities [12][13]	No single federal-level training program; advanced through the *Digital Education Action Plan 2021–2027* and working groups led by the European Digital Education Hub (EDEH) [8]
Policy instruments and investment	The Action Plan was jointly issued by the Ministry of Education and multiple other ministries, spanning basic education through lifelong learning, with a target of essentially completing an integrated AI education system by 2030 [6][7]	Relies mainly on existing discretionary grants, public–private partnerships, and the Presidential AI Challenge; a White House task force on AI in education has been established [12][13]	Constrained by the compliance framework of the EU AI Act (2024), embedding teacher competency within the broader digital-literacy and compliance regime [9]
Training coverage (primary data, 2024–2026)	2.97 million teachers trained cumulatively; 700+ AI tools developed for teaching, research, and professional learning (as of end-2024) [10][11]	43% of teachers report having received AI training at least once (fall 2024); 48% of districts have run teacher AI training, up sharply from 23% in fall 2023 [2][3]	[TBD: official EU-level data on teacher AI training coverage]
Primary implementing bodies	The Ministry of Education and co-issuing ministries; local education authorities and the National Smart Education Platform [6][7]	The federal Department of Education, the National Science Foundation (NSF), the White House AI education task force; states and districts retain substantial autonomy [12][13]	The European Commission's Directorate-General for Education and Culture, the European Digital Education Hub (EDEH), and competent authorities in member states [8][9]

Note: China's "AI+ Education" Action Plan was jointly issued by the Ministry of Education and relevant departments in April 2026 [6][7]. U.S. Executive Order 14277, "Advancing Artificial Intelligence Education for American Youth," was signed on April 23, 2025 [12][13]. The EU's Ethical Guidelines were first published in 2022 and revised alongside the 2024 EU AI Act [8][9]. Any specific figures or state-level details in this table that could not be verified against an official primary source are marked [TBD] rather than filled in with an estimate.

9.3.2 A Supranational Public Benchmark

At the supranational level, UNESCO's September 2024 *AI Competency Framework for Teachers* offers a globally reference-worthy set of consensus dimensions (see 9.2.1 and sources [1][4]), against which national standards can be compared. UNESCO itself positions the framework as a tool "to guide countries in developing their own competency frameworks, training programs, and assessment parameters" rather than as an enforceable standard [1]. One caveat: if the framework's potential reach — number of countries covered, teachers reached — is ever cited, it should be explicitly flagged as third-party estimation rather than official statistics, to avoid misreading. This chapter does not make unverified claims about scale.

9.3.3 A Qualitative Comparison of Three Institutional Approaches

Taken together, the primary sources point to three distinct orientations:

- **China: top-down, unified deployment with system-wide embedding.** The defining feature is writing teacher AI literacy directly into **licensure and certification** — the most binding institutional lever available. "Incorporation into teacher qualification exams" turns literacy from an aspiration into an access requirement [6][7]. Paired with large-scale training through the national platform (on the order of 2.97 million teacher-trainings) [10][11], this forms a closed "standard–assessment–training" loop. The risk lies in the tension between a unified national standard and considerable local variation.
- **United States: federal direction, with states and districts retaining substantial autonomy — a fragmented landscape.** The federal government sets direction through executive order and uses grant priority as leverage (120-day priority for teacher training) [12][13], but does not set a unified competency standard; specific standards and training are left to states and districts. The upside is responsiveness to local context and flexibility to iterate; the cost is highly uneven coverage — RAND's data show a training-provision gap of 67% in low-poverty districts versus 39% in high-poverty districts [2][3].
- **European Union: competency embedded within a unified regulatory framework.** Rather than standing up a separate teacher AI standard, the EU folds AI competency into the existing DigCompEdu digital-literacy framework, and governs teacher compliance jointly through the Ethical Guidelines and the binding EU AI Act [8][9]. The underlying logic is "compliance first,

competency embedded" — strong on ethics and data protection, but comparatively opaque on training rollout data at scale (see [TBD] entries in the table above).

Worth adding is the United Kingdom's distinctive path. The UK Department for Education (DfE) published an updated *Generative Artificial Intelligence (AI) in Education* guidance in June 2025, and — working with the Chartered College of Teaching and the Chiltern Learning Trust — provides supporting materials that include training modules, audit tools, and a **free certified assessment**. Teachers who complete the training and pass the online assessment earn credentials attesting to "safe, ethical AI use" [14]. This "guidance–training–micro-credential" package is a real-world national instance of the micro-credentialing model discussed in Section 9.4.4.

These three (or four) paths each follow their own institutional logic for enabling infrastructure, and no simple ranking applies. But one common trend stands out: across 2024–2026, every jurisdiction examined here is moving from "encouraging use" toward "institutionalized competency-building" — the difference lies only in which lever does the institutionalizing (qualification exams, grant leverage, compliance frameworks, or micro-credentials).

9.4 Models for Delivering Professional Development, and Their Effectiveness

Once we know what literacy teachers need, the next question for enabling-infrastructure design is how to build it effectively. This section surveys the mainstream models currently used for teacher AI professional development and offers a qualitative judgment of where each fits.

9.4.1 Centralized Training

Short workshops and topical lectures are the default vehicle for fast, broad-reaching rollout, since they are easy to scale and cheap to organize. This is also, in practice, the largest model by volume today: China's 2.97 million teacher-trainings, and the jump in U.S. district-level training provision from 23% in fall 2023 to 48% in fall 2024 [10][11][2][3], both fall into this category. But centralized training runs into a familiar problem everywhere: training disconnected from practice, and one-off input that fails to stick. RAND's own observations bear this out — most training still sits at a **very early stage**, aimed at making teachers "comfortable" with AI tools, and its content responds more to teachers' anxieties about AI than it builds systematic competency [2]. Centralized training, in short, works well for quickly spreading foundational understanding, but has limited reach into more practice-dependent dimensions like pedagogical integration and prompt engineering.

9.4.2 School-Based Professional Learning and Communities of Practice

This model embeds professional development into teachers' everyday collaborative work — peer collaboration and lesson refinement driving continuous improvement. It carries forward the emphasis the original master white paper placed on online teacher communities, and has found a new vehicle in the generative-AI era: teachers building reusable practice assets around prompt design and AI-assisted lesson preparation. Recent academic reviews of generative-AI teacher professional development likewise conclude that effective integration of AI teaching tools requires **systematic training paired with ongoing technical support**, not one-off input — and that "how to design effective professional development for pre-service and in-service teachers alike" remains a core open question in the field [15][16]. The limitation here is a heavy dependence on a school's existing culture of collaborative teaching research.

9.4.3 Agent-Assisted Personalized Professional Learning

An emerging model that uses AI itself as the support tool for teacher professional development — offering personalized learning paths, real-time feedback, and coaching-style rehearsal. It follows the same logic found in this institute's research on educational robotics and AI glasses: using AI to empower the professional. UNESCO's framework lends this approach international legitimacy directly, listing "AI for professional learning" as one of its five dimensions [1][4]. On the supply side, this model already has scale behind it: alongside its teacher-training push, China has developed more than 700 AI tools supporting teaching, research, and professional learning [10][11]. But the **effectiveness** of this model at scale still lacks robust evidence — existing reviews are largely descriptive, framed around "potential" and "perceived usefulness," with long-horizon, controlled effectiveness evidence still thin [15][16]. This Blue Book recommends approaching it with cautious optimism.

9.4.4 Micro-Credentials and Competency Portfolios

This model documents teachers' professional growth through granular, verifiable competency credentials, mapping naturally onto the four-dimensional structure described above. Its value lies in converting "attended training" into "demonstrated competency" — a verifiable credential rather than a mere participation record. The DfE's free certified assessment, delivered jointly with the Chartered College of Teaching, is a national-scale real-world instance of this model: teachers complete DfE training materials and then sit an online assessment, **producing evidence** of their understanding of safe, ethical AI use [14]. China's approach of "incorporation into teacher qualification exams" can likewise be read as a mandatory variant of the same micro-credentialing logic [6][7]. Open questions for this model include the credibility and cross-regional recognition of certification standards, and how to avoid assessment drifting into "teaching to the test" — divorced from actual classroom competency.

9.4.5 Where the Effectiveness Evidence Stands, and This Institute's Position

The current state of academic evidence on these models' effectiveness can be summarized in one line: **strong directional consensus, weak causal evidence.** Systematic reviews consistently confirm generative AI's positive potential for teaching performance — improved usability, personalized resources, automated feedback — and consistently stress that teachers need systematic, sustained professional development [15][16]. But on the more decisive question of which specific professional-development approach works best, high-quality evidence from long-horizon, tightly controlled studies remains scarce [15][16]. Given this, this institute holds a reserved view toward any single model claiming "best practice" status, and instead advocates a context-sensitive combination: centralized training to quickly establish foundational understanding, school-based professional learning and communities of practice to consolidate the practice-dependent dimensions, micro-credentials to provide accountable proof of competency, and a disciplined "evaluate first, scale later" posture toward agent-assisted professional learning.

9.5 The Teacher Literacy Gap Through an Equity Lens

Teacher AI literacy is not evenly distributed, and its unevenness tends to compound existing gaps in educational resources rather than stand apart from them. This section connects to the dedicated equity-of-access discussion in Chapter 6, focusing here on the structural gap on the teacher side and anchoring its scale in primary data.

Urban–rural, regional, and wealth gaps. This gap is already backed by clear quantitative evidence. RAND's fall 2024 survey found that the share of U.S. low-poverty districts providing AI training to teachers (67%) significantly exceeds the share for high-poverty districts (39%) — a 28-percentage-point gap [2][3]. And based on district-level planning projections, this gap is not expected to narrow by fall 2025 — it may widen further: nearly all low-poverty districts will have completed teacher training, while only about 60% of high-poverty districts are expected to do so [2][3]. The danger here is structural: AI was supposed to help close educational gaps, but if the teacher literacy gap goes unaddressed, the technology risks amplifying disparities instead. In China's case, the 2.97-million-teacher training push reflects a broad, inclusive intent at the aggregate level [10][11], but differences in training quality and tool accessibility across urban and rural areas and across regions still need to be monitored with more granular data ([TBD: official disaggregated data comparing urban and rural teacher AI training quality in China]).

Generational and subject-area gaps. Within the teaching profession itself, acceptance of and facility with AI tools vary considerably by age, subject taught, and baseline digital literacy. The very design of UNESCO's "Acquire–Deepen–Create" progression implicitly acknowledges that teachers start from different points — a uniform proficiency requirement, absent a tiered pathway to reach it, risks compounding the frustration and attrition of teachers who start further behind [1][4]. The finergrained distribution of this gap within the teaching workforce still needs to be filled in with local survey data ([TBD: quantitative data on generational/subject-area differences within the teaching workforce]).

Special education and vulnerable student populations. Teachers serving students with special needs face additional professional challenges in adapting general-purpose AI tools to specific support contexts, and targeted training resources for this group remain broadly scarce ([TBD: evidence/sources on AI training resources for special-education teachers]). This points to a design principle for equity in the enabling infrastructure: it cannot stop at "getting more teachers using AI" — it must also ask whether teachers in different circumstances can use AI well enough to actually serve their students. UNESCO's decision to list "human-centredness" as its lead dimension, and the EU's use of ethical guidelines to constrain AI use [1][8], both point toward the same conclusion: equity is not a side issue for training — it belongs inside the competency framework itself, as a built-in constraint.

9.6 Conclusions and Recommendations

This chapter's central judgment is this: among all the factors shaping the global K-12 AI education enabling infrastructure, teacher AI literacy carries the dual distinction of being both the **highest-leverage** factor and the **most widespread weak point**. Two sets of primary evidence support this. First, UNESCO has already confirmed at the international level that "most countries have yet to define teacher AI competencies, and most educators lack guidance" [1]. Second, even in the United States, where training has expanded rapidly, more than half of teachers had still received no AI training as of fall 2024, and training coverage shows a marked wealth gap [2][3]. The return on investment in policy, platforms, and resources depends heavily on whether teachers can convert that investment into classroom learning gains — and the current overall readiness of the teaching workforce is not yet sufficient to support generative AI and agentic classrooms at scale.

At the same time, policy practice from 2024–2026 offers instructive positive examples. China has written AI literacy into teacher qualification exams and delivered 2.97 million teacher-trainings through its national platform [6][7][10][11]. The United States has used executive order and grant leverage to mobilize teacher training within 120 days [12][13]. The EU has embedded competency within DigCompEdu and its ethical guidelines, folded into a compliance framework [8][9]. The UK has advanced micro-credentialing through its three-part "guidance–training–free certification" package [14]. These levers differ, but together they confirm that this chapter's directional recommendations rest on real institutional footing.

Based on the analysis above, this institute puts forward the following directional recommendations (deliberately without specific quantitative targets, which each jurisdiction should set according to its own circumstances):

1. **Establish a teacher AI competency standard aligned with the four-dimensional structure above and mappable onto UNESCO's five-dimensional framework**, and progressively fold it into teacher professional-development and certification systems — drawing on China's "incorporation into qualification exams" and the UK's "free certified assessment" as institutionalization models — so that "literacy" moves from an aspirational concept to a measurable, accountable institutional requirement [1][6][14].
2. **Shift the center of gravity from one-off centralized training to sustained professional development embedded in everyday teaching practice.** Existing evidence shows that centralized training mostly stalls at the stage of "making teachers comfortable" [2]; teacher communities of practice need to be rebuilt and upgraded so that newer competencies — prompt engineering, human–AI collaboration — take root in real classroom practice [15][16].
3. **Treat equity as a precondition for teacher professional development, not an afterthought.** Given that training coverage already shows a clear wealth gap (67% versus 39%) [2][3], resource allocation should structurally favor rural areas, less-developed regions, and special-education teachers, to prevent the technology from amplifying existing gaps.
4. **Introduce agent-assisted teacher professional learning cautiously.** Even as "AI for professional learning" has UNESCO's endorsement and the supply side already has tools at scale (700+) [1][10], the reality of thin long-horizon effectiveness evidence must be taken seriously [15][16]; a localized evidence-based evaluation mechanism should be built before scaling up, to avoid low-value investment dressed up as "using AI to train AI users."
5. **Strengthen teachers' ethical and data literacy as the last link in the governance chain.** Echoing UNESCO's elevation of "ethics" and "human-centredness" to independent dimensions, and the EU's approach of constraining teacher behavior through ethical guidelines and the AI Act [1][8][9], this recommendation connects to the governance framework in Chapter 10, translating the principles of academic integrity, data protection, and algorithmic fairness into judgment calls teachers can actually apply day to day.

The teacher-literacy baseline this chapter establishes is a precondition for understanding whether the governance framework in Chapter 10 can actually be implemented. If governance rules lack teachers equipped to enforce them, the institutional design remains suspended at the level of text.

Chapter References

[1] AI competency framework for teachers · UNESCO · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-teachers

[2] More Teachers Than Ever Before Are Trained on AI. Are They Ready to Use It? · Education Week · 2025 · https://www.edweek.org/technology/more-teachers-than-ever-before-are-trained-on-ai-are-they-ready-to-use-it/2025/04

[3] More Districts Are Training Teachers on Artificial Intelligence: Findings from the American School District Panel · RAND Corporation · 2025 · https://www.rand.org/pubs/research_reports/RRA956-31.html

[4] What you need to know about UNESCO's new AI competency frameworks for students and teachers · UNESCO · 2024 · https://www.unesco.org/en/articles/what-you-need-know-about-unescos-new-ai-competency-frameworks-students-and-teachers

[5] UNESCO's ICT Competency Framework for Teachers · UNESCO · 2023 · https://www.unesco.org/en/digital-competencies-skills/ict-cft

[6] China launches AI empowering education action plan, includes AI into teacher qualification exams · Global Times · 2026 · https://www.globaltimes.cn/page/202604/1358611.shtml

[7] China aims to build an AI literacy system · The State Council of the People's Republic of China (english.www.gov.cn) · 2026 · https://english.www.gov.cn/news/202604/15/content_WS69df29e6c6d00ca5f9a0a6b1.html

[8] Ethical guidelines on the use of artificial intelligence and data in teaching and learning for educators · European Commission, European Education Area · 2022 · https://education.ec.europa.eu/node/2285

[9] Guidelines on the ethical use of artificial intelligence and data in teaching and learning for educators · Publications Office of the European Union · 2025 · https://op.europa.eu/en/publication-detail/-/publication/f692aa0b-17a7-11f1-8870-01aa75ed71a1/

[10] Chinese schools embrace AI for smarter learning · People's Daily Online · 2025 · https://en.people.cn/n3/2025/0722/c90000-20343341.html

[11] China advances AI curriculum to cover full basic education · State Council Information Office (english.scio.gov.cn) · 2025 · http://english.scio.gov.cn/pressroom/2025-05/13/content_117871666.html

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

[13] Executive Order: Advancing Artificial Intelligence Education for American Youth · Holland & Knight (Insights) · 2025 · https://www.hklaw.com/en/insights/publications/2025/04/executive-order-advancing-ai-education-for-american-youth

[14] Generative artificial intelligence (AI) in education · UK Department for Education (GOV.UK) · 2025 · https://www.gov.uk/government/publications/generative-artificial-intelligence-in-education/generative-artificial-intelligence-ai-in-education

[15] Reimagining teacher development in the era of generative AI: A scoping review · Teaching and Teacher Education
(ScienceDirect) · 2025 · https://www.sciencedirect.com/science/article/abs/pii/S0742051X25003130

[16] Artificial intelligence in teaching and teacher professional development: A systematic review · Computers and Education: Artificial Intelligence
(ScienceDirect) · 2024 · https://www.sciencedirect.com/science/article/pii/S2666920X24001589

Chapter 10 Equity of Access, Safety and Ethics, and Governance; Conclusion and Implementation Roadmap

Literacy, curriculum, human–AI collaboration, and assessment together define what to teach, how to teach it, and how to evaluate it. But whether these translate into genuine opportunity for every learner depends on the last — and most fundamental — line of defense: whether the enabling infrastructure is equitably accessible, safe and trustworthy, and effectively governed. This chapter draws the book's capability-focused threads together into a governance lens, asking who is being left behind, who is bearing the risk, and who is accountable — and then lays out a phased implementation roadmap for 2027 and beyond. The chapter's central judgment is this: across 2024–2026, as generative and agentic technologies have spread rapidly, deployment across major economies has systematically outpaced the buildout of governance and safeguards. As a result, the equity–safety–governance triad is no longer a side issue bolted onto the main agenda — it is fast becoming the **dominant variable** determining whether AI's net effect on education is positive or negative.

10.1 Equity of Access: From the Access Divide to the Capability Divide

The arrival of generative AI in primary and secondary schools has structurally reshaped what we mean by the digital divide. In the pre-generative era, the divide was mainly about **access to devices and connectivity (the first-level divide)**. Between 2024 and 2026, as foundation models became reachable through ordinary browsers and mobile devices, the fault line shifted along two new axes. The first is a **capability divide in use and literacy (the second-level divide)**: even when access is equal, different groups differ sharply in whether they use AI for deep learning or for shallow shortcuts, and in whether they can wield the tools critically. The second is an **emerging AI divide in compute, quality models, and data**: premium paid models, stable computing capacity, and localized educational data remain unevenly distributed across regions and schools.

It bears emphasizing that the access divide has by no means disappeared globally — its center of gravity has simply shifted from "is there a connection at all" to "is the connection good enough, and can people afford it." According to the International Telecommunication Union's (ITU) *Facts and Figures 2024*, roughly 2.6 billion people worldwide remained offline in 2024, or 32% of the global population; about 1.8 billion of the offline population lives in rural areas [1]. The urban–rural gap is stark: global internet use stands at roughly 83% in urban areas versus under half (about 48%) in rural areas [1]. The income divide is equally glaring — internet use is about 93% in high-income countries versus only about 27% in low-income countries, where a fixed broadband subscription can cost roughly a third of average monthly income [1]. In other words, while developed-country discourse has already moved on to "AI literacy," a substantial share of the world's school-age children remain stuck at the access layer. This Blue Book therefore argues for diagnosing equity through **three cumulative layers rather than one**: the access divide and the capability divide are not sequential stages where one replaces the other — they coexist and compound, differently, from region to region.

10.1.1 A Three-Layer Framework for Access

To keep "equity" from collapsing into a slogan, this Blue Book organizes both diagnosis and response around a three-layer access framework:

- **Access layer**: the availability and stability of devices, network bandwidth, and model services.
- **Use layer**: language support, accessibility, teacher mediation, and curricular time — the factors that determine whether "having access" actually turns into using AI well, often, and effectively.
- **Outcome layer**: whether AI narrows or widens existing gaps in academic achievement and development — the ultimate test of equity.

The policy value of this three-layer framework is that it turns "equity" from a moral slogan into something that can be measured, and held accountable, layer by layer. The OECD's analysis of PISA 2022 provides cross-national evidence for exactly this pattern: schools with higher socioeconomic status are generally better at using digital devices to strengthen teaching and learning than disadvantaged schools are, and advantaged students report better access to quality ICT resources than disadvantaged students do [2]. In other words, achieving universal device rollout at the access layer does not automatically produce equity at the use and outcome layers — it can instead **widen** existing gaps, because advantaged groups are simply better positioned to exploit the tools. This is the Matthew effect taking concrete shape in AI-enabled education.

Access tier	Key indicators	Verifiable evidence
Access layer	Connectivity rate, bandwidth, availability of model services	2.6 billion people offline globally (32%); rural 48% vs. urban 83%; low-income countries 27% vs. high-income countries 93% [1]. China: urban and rural school internet coverage and interactive-multimedia equipment rates both exceed 90% [3]
Use layer	Teacher AI literacy, accessibility adaptation, native-language/low-resource-language support, distribution of quality resources	OECD: advantaged schools outperform disadvantaged schools at using digital devices for teaching; ICT resource access shows a socioeconomic gap [2]. China: urban–rural gaps remain substantial in hardware adequacy, technical support, and capacity-building training [3]
Outcome layer	Correlation between AI use and	OECD: students who use devices

	academic value-added; changes in between-group gaps	moderately for learning perform better and report a stronger sense of belonging; students frequently distracted by digital devices score markedly lower in mathematics [4] (see 10.2.2)

10.1.2 Urban–Rural and Regional Gaps

In the generative-AI era, the urban–rural gap has taken on a distinctive shape: devices are converging while capability is diverging. Terminals and connectivity keep improving under national targeted-support programs, but the ability to pay for premium models, access to stable compute, instructional-research support, and teacher literacy remain heavily concentrated in major cities and economically developed regions.

China illustrates how solid the "access layer" foundation can become. According to a Ministry of Education review, by the relevant reporting date, internet access across the country's primary and secondary schools (including teaching sites) had reached near-universal coverage, with urban and rural internet coverage and interactive-multimedia equipment rates both above 90% [3]. The National Smart Education Platform for Primary and Secondary Schools has become one of the largest public-service systems for educational resources in the world: by the end of 2023, its cumulative registered users had surpassed 100 million, with more than 36.7 billion page views; by May 15, 2024, total page views had reached 40.54 billion, drawing together roughly 88,000 K-12 resources [5]. The significance of this kind of national public platform lies precisely in its **equity function** — by delivering quality curricular resources to rural and under-resourced schools on a "state-funded, universally free" basis, it offsets the tendency of market-driven quality resources to concentrate in developed regions.

Yet the "capability divergence" beneath this "access convergence" cannot be overlooked. A nationwide survey by Peking University's China Institute for Educational Financial Research found that urban and rural schools differ relatively little in network environment, software, and resources, but differ substantially in **the adequacy of school-provided hardware, technical support, and training to build teachers' information-technology capacity** [3]. In other words, the real gap today is no longer "who has the resources" but "who has people who know how to use them, and who has the support to use them well." Measuring equity by access rate alone would systematically mask a widening hidden gap at the use and outcome layers — which is exactly the evidence base behind this Blue Book's repeated emphasis on moving equity upstream, into capacity-building and procurement admission criteria.

10.1.3 Special-Needs Populations and Accessibility

Generative AI carries inherent accessibility potential — real-time captioning, speech-to-text, text simplification, cross-modal conversion, and adjustable difficulty levels can, in principle, deliver low-cost "reasonable accommodation" to learners who are blind or have low vision, are deaf or hard of hearing, have learning disabilities, are on the autism spectrum, or otherwise have special educational needs (SEN). This potential aligns closely with international consensus: UNESCO's *AI Competency Framework for Students*, released in September 2024, places "human-centered values, ethics, sustainability, and inclusion" as a value thread running through the entire framework, explicitly requiring that students engage with AI safely and meaningfully [6][7].

But realizing this potential depends on preconditions that are themselves sources of equity risk. First, models can misread minority accents, non-standard speech patterns, sign language, and augmentative-and-alternative-communication methods — meaning a tool billed as serving everyone can, in practice, re-exclude SEN learners. Second, it matters whether accessibility features are **on by default** rather than buried in settings menus requiring extra configuration. Third, interoperability between assistive technologies and instructional systems remains a real constraint. On this basis, this Blue Book argues that **accessibility by design should be a default requirement — not an optional feature — in the procurement and admission of educational AI**: products should be required, at the point of admission, to demonstrate that accessibility capabilities are enabled by default and to publish their adaptation and evaluation results for typical SEN scenarios, rather than shifting the burden of proving "reasonable accommodation" onto the most vulnerable users. A concrete tool inventory and evidence-based outcome data for SEN learners appear at [TBD: data/source]; this Blue Book makes no quantitative claims where reliable primary evidence is unavailable.

10.1.4 Language Equity and Localization

Most mainstream foundation models perform best in high-resource languages, while their capability, safety, and cultural fit are noticeably weaker in minority, dialect, and low-resource languages — an underappreciated equity risk. Language equity is not simply about whether students can ask questions in their native language; it also concerns whether generated content carries hidden cultural bias, whether it fits local curricular context, and whether it might inadvertently erode intergenerational transmission of language and culture. UNESCO's competency framework places "AI ethics" and "inclusion" at its core [6][7], and the underlying logic points directly here: an AI that serves only high-resource-language users is, by construction, not inclusive. Data on multilingual model evaluation and the state of localized educational-corpus building appear at [TBD: data/source]. This Blue Book recommends making "usability and safety in native and low-resource languages" an explicit item in admission evaluations, and encourages vendors, research institutions, and local education authorities to collaborate on building localized, debiased educational corpora.

10.1.5 Takeaway: Write Equity into Admission Criteria

The governance implication of the access findings is clear: **equity cannot rely on after-the-fact compensation alone — it must be built upstream as a hard constraint on design and procurement.** The access gap documented by the ITU [1], the use and outcome gaps documented by the OECD [2][4], and the structural risks around accessibility and language together point to the same conclusion: left to the market's own devices, AI is more likely to amplify gaps than to close them. This judgment is translated into concrete governance mechanisms in Section 10.3 and into a roadmap in Section 10.5.

10.2 Safety, Ethics, and the Risk Landscape

Education has far less tolerance for safety and ethical missteps than general consumer contexts do, because its subjects are minors, its stakes are long-term development, and the power relationship is inherently unequal. When UNESCO pushed for the first global guidance in this space, it flagged the urgency with a single stark data point: in a global survey of more than 450 schools and universities, fewer than 10% of institutions had adopted formal institutional policies or guidance on the use of generative AI applications — a gap driven "in large part" by the absence of regulation at the national level [8]. In practical terms, this means the vast majority of minors are encountering powerful generative tools in an environment with almost no institutional guardrails. This section maps that risk landscape systematically, without dwelling on individual incidents.

10.2.1 A Risk Taxonomy for Minors

- **Content-safety risk**: generation or exposure to harmful, age-inappropriate, violent, or manipulative content.
- **Data and privacy risk**: excessive collection, secondary use, and cross-border transfer of minors' personal information, learning-behavior data, and biometric data.
- **Psychological and relational risk**: emotional dependence induced by anthropomorphic interaction design, quasi-familial or quasi-romantic companionship dynamics, and the displacement of peer and teacher–student relationships.
- **Cognitive and academic-integrity risk**: outsourced-thinking use patterns that erode independent thought, excessive "cognitive offloading," and blurred boundaries around academic misconduct (discussed further in Chapters 7 and 8, on human–AI collaboration and assessment).
- **Equity and bias risk**: model bias amplified in high-stakes contexts such as grading, tracking, and recommendation.
- **Reliability risk**: hallucinations and factual errors propagating through teaching and assessment.

· **Attention risk**: sustained competition from devices and applications for minors' attention (see the PISA evidence below).

10.2.2 Mechanisms and Exposure Points of Representative Risks

Risk category	Triggering mechanism	Primary points of exposure	Affected populations
Content safety	Unpredictable generation, jailbreaks that bypass safeguards	Open-ended dialogue, homework help	All students, especially younger children
Data privacy	Default connectivity, log retention, third-party sharing, biometric collection	Registration, interaction, assessment	All students and teachers
Emotional dependence	Anthropomorphic design, constant availability, infinite patience	Companion/study-buddy applications	Minors with higher emotional needs
Academic integrity	High-quality generation that is hard to trace	Homework, exams, inquiry-based work	All students
Bias and discrimination	Bias embedded in training data	Automated grading, ability tracking	Disadvantaged and minority groups
Attention competition	Notification/recommendation mechanisms, task-switching	Device use in and out of class	All students

Among these risks, "attention and focus" stands out as one of the few areas already backed by large-sample, cross-national quantitative evidence, making it a useful case study in evidence-based governance. According to OECD PISA 2022 (Volume II), an average of 65% of students across OECD countries reported being distracted by digital devices in "at least some" math classes, and about 30% reported being distracted in "most or every" math class [4]. Meanwhile, roughly 45% of students reported "feeling anxious" when their phone was not nearby [4]. More critically, this connects directly to academic outcomes: students who reported being distracted in most or every class scored roughly 15 points lower in mathematics than students who reported almost never being distracted — and a full year of learning progress for a 15-year-old is worth roughly 20 points on the same scale [4]. The policy implication cuts both ways rather than pointing in a single direction — PISA also found that 15-year-olds who used digital devices **moderately** for schoolwork tended to perform better and report a stronger sense of belonging at school [4]. The goal of safety governance, then, is not to ban technology outright, but to steer use toward the "moderate, purposeful, teacher-mediated" zone. As for the measured incidence rates and documented case evidence for risks such as content safety, jailbreaking, and emotional dependence, no recognized reliable primary statistics currently exist; see [TBD: data/source] — this Blue Book does not attempt inferential quantification here.

10.2.3 Translating Ethical Principles into Educational Practice

UNESCO and the regulatory frameworks of major economies have converged on a considerable degree of consensus around common ethical principles. The OECD Council's *Recommendation on Artificial Intelligence* (adopted 2019, revised 2024 — the first intergovernmental AI standard, covering its 38 member countries and beyond) establishes five value-based principles: inclusive growth, sustainable development, and well-being; respect for the rule of law, human rights, and democratic values, including fairness and privacy; transparency and explainability; robustness, security, and safety; and accountability [9]. UNESCO's 2023 *Guidance for Generative AI in Education and Research* (authored by Fengchun Miao and Wayne Holmes, published September 2023) takes a "human-centered" approach as its overarching principle, recommending to its 193 member states a set of key regulatory steps including data-privacy protection and a minimum age threshold for independent interaction with generative AI platforms — the guidance suggests age 13 as a reference minimum for independent classroom use of such tools [8][10].

This Blue Book emphasizes that these principles carry little force if they remain declarations — they need to be **translated into education-specific, operable rules**. Internationally, legislation has already begun moving from "principle" to "hard rule," most notably in the EU's *Artificial Intelligence Act* (EU AI Act): it classifies multiple categories of educational AI systems — automated grading, admissions screening, exam proctoring, student monitoring — as **high-risk** (Annex III), requiring quality-management systems, risk assessment, and human oversight; public educational institutions must conduct a Fundamental Rights Impact Assessment (FRIA) under Article 27 before deploying high-risk systems [11]. Particularly significant is Article 5(1)(f), which **expressly prohibits** AI systems that infer the emotions of natural persons from biometric data in educational institutions and workplaces (with exceptions for medical and safety purposes) — grounded in concerns about the questionable scientific validity of such systems and the inherent power imbalance in educational settings [11]. This provides a binding international precedent for defaulting to a ban on emotion recognition and profiling-style assessment for minors.

Building on this, this Blue Book proposes translating these general principles into a set of checkable, education-specific rules: age-tiered use with parental awareness for minors; a default ban on emotion recognition and profiling-style assessment aimed at minors; conspicuous disclosure of identity in anthropomorphic interactions (so minors are not misled into believing an AI is a real person or "friend"); and minimization of educational-data collection with a preference for local processing. Specific provisions already in force across relevant jurisdictions are cross-referenced in Section 10.3 and the chapter references.

10.3 Governance Mechanisms: A Multi-Level, Multi-Stakeholder Framework

Delivering on both safety and equity depends on governance, not on any single technical fix. This Blue Book proposes a governance framework spanning three levels — macro, meso, and micro — and three phases — ex ante, in-process, and ex post — anchored in real legislation and policy documents from 2024–2026 across multiple jurisdictions.

10.3.1 Three Levels of Governance Actors

- **Macro level (policy and standards)**: national/regional laws and regulations, curriculum standards, admission and safety norms.
- **Meso level (regional authorities and schools)**: procurement admission by education authorities and schools, codes of use, teacher capacity-building, and incident response.
- **Micro level (classroom and family)**: teachers' mediating responsibility, students' responsible use, and parents' awareness and involvement.

The macro level has entered a phase of intensive lawmaking. Beyond the OECD Recommendation [9] and the UNESCO guidance and competency frameworks [8][10][6][7] and the EU AI Act [11] already discussed, the Council of Europe's *Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law* — opened for signature in Vilnius on September 5, 2024 — became the **world's first legally binding international AI treaty**. Drafted by the Council of Europe's 46 member states with participation from observers including the United States, Canada, Japan, and the EU, along with several non-member states, it establishes core principles of transparency, accountability, non-discrimination, and human-rights protection [12]. This marks AI governance's shift from soft-law initiative to hard-law obligation.

China, meanwhile, has built a combination of "curriculum guidance + use norms + minor-protection legislation." The Ministry of Education's Teaching Guidance Committee for Basic Education released the *Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition)* and the *Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition)* on May 13, 2025, adopting a spiral progression across primary, middle, and high school and setting requirements for stage-appropriate generative-AI use and data security — for instance, requiring primary-school students to "build a basic understanding of privacy protection and digital identity," and requiring administrators to "establish AI education data-security management norms ... and build privacy-protection mechanisms" [13]. On minors' online and data protection, China's institutional foundation includes State Council Order No. 766, the *Regulation on the Protection of Minors in Cyberspace* (adopted at a State Council executive meeting on September 20, 2023; in effect since January 1, 2024) [14]; along with rules established under the *Personal Information

Protection Law* and the *Law on the Protection of Minors (as amended in 2024)*, which require that processing the personal information of minors under 14 obtain consent from a parent or other guardian, and prohibit compelling consent to non-essential personal-information processing [15].

The United States has taken an approach that leans more toward "guidance plus overlay on existing privacy law." The U.S. Department of Education released *Designing for Education with Artificial Intelligence*, a guide for developers, in July 2024, followed by a *Toolkit for Safe, Ethical, and Equitable AI Integration* for schools in October 2024, covering risk mitigation, student-privacy protection, and digital equity, among other modules [16]. On the data-protection side, the Federal Trade Commission (FTC) finalized amendments to the Children's Online Privacy Protection Rule (COPPA Rule) in early 2025, bringing "biometric identifiers (such as voiceprints and facial characteristics)" within the scope of protected "personal information" and requiring separate opt-in consent for sharing children's data with third parties; the amendment took effect June 23, 2025, with a compliance transition period for covered entities running roughly through April 22, 2026 [17]. This amendment carries direct implications for educational AI systems that widely process voice and image data.

10.3.2 Comparing Governance Pathways Across Major Jurisdictions

The table below structures the primary-source documents from 2024–2026 verified in this chapter, to illustrate a governance spectrum in which the same issue is addressed through different pathways.

Jurisdiction/Body	**Landmark document (year · reference/nature)**	**Legal force**	**Key K-12 constraint**
UNESCO	*Guidance for Generative AI in Education and Research* (2023); *Teacher/Student AI Competency Frameworks* (2024) [8][10][6][7]	Soft law/recommendation	Suggested age-13 minimum; data-privacy protection; ethics and inclusion built into competency frameworks
OECD	*Recommendation on Artificial Intelligence* (adopted 2019, revised 2024) [9]	Intergovernmental standard (soft law)	Five value-based principles (including fairness, privacy, transparency, accountability)
European Union	*Artificial Intelligence Act* (EU AI Act, 2024) [11]	Hard law (regulation)	Multiple education scenarios classified as high-risk; school-based emotion recognition banned; FRIA required for high-risk systems
Council of Europe	Framework Convention on AI and Human Rights (2024,	Hard law (international treaty)	First binding international AI treaty; transparency,

	signed in Vilnius) [12]		accountability, non-discrimination
China	Two K-12 AI education/generative-AI use guidelines (2025) [13]; *Regulation on the Protection of Minors in Cyberspace* (State Council Order No. 766, effective 2024) [14]; PIPL/Law on the Protection of Minors [15]	Guidance (soft law) + administrative regulation/statute (hard law)	Stage-appropriate use norms; education data-security norms; guardian consent required for children under 14
United States	Department of Education's *Designing for Education with AI* (2024, developer guide) and AI toolkit (2024) [16]; FTC COPPA amendments (effective 2025) [17]	Guidance (soft law) + federal rule (hard law)	Privacy/equity/safety guidance; biometric data added to children's personal information, requiring separate consent

Three patterns emerge from this comparison. First, **a combination of soft and hard law, with hard law providing the backstop** — UNESCO and the OECD supply value consensus, while the EU and the Council of Europe harden it into regulations and treaties, and China and the United States implement it through a combination of guidance and existing law. Second, **special protection for minors' data is a cross-jurisdictional consensus** — whether through COPPA's inclusion of biometric data or China's guardian-consent requirement for children under 14, all paths converge on stricter regulation of minors' biometric and identity data. Third, **"emotion recognition/profiling" is emerging as a red line in education** — the EU has already legislated a ban, and other jurisdictions' "default-disabled" recommendations echo the same direction. These three patterns form the legislative basis for this Blue Book's procurement admission checklist.

10.3.3 Governance Across the Full Lifecycle

Phase	Key actions	Responsible actors	Tools/mechanisms
Ex ante (admission)	Safety and equity impact assessment, age-appropriateness review, procurement admission checklist	Education authorities, schools	Fundamental rights/safety impact assessment (per EU AI Act Article 27 FRIA [11]), admission checklist
In-process (operation)	Codes of use, teacher oversight, logging and data minimization, explainable	Schools, teachers, vendors	Codes of use, audit logs, data minimization (per the Minor Protection

	feedback		Regulation/PIPL [14][15])
Ex post (accountability)	Incident reporting, appeals and remediation, outcome and equity review	All stakeholders	Appeals mechanism, annual assessment

10.3.4 A Procurement Admission Checklist for Educational AI (Proposed Framework)

This Blue Book proposes translating abstract principles into checkable items through a checklist, for use by education authorities and schools at the point of admission. Each item traces back to the real legislation discussed above (specific thresholds and compliance items should be adapted to locally applicable law):

· **Age-appropriateness and content safety**: age tiering, content filtering, and human backstop in place; alignment with reference thresholds such as the age-13 minimum [8][10].

· **Data and privacy**: data minimization, retention limits, and controllable cross-border/third-party sharing; guardian consent implemented for users under 14 [15]; separate consent obtained for biometric data (voiceprint/facial) [17].

· **Transparency and explainability**: disclosed model identity (conspicuous labeling in anthropomorphic interactions), clearly stated capability boundaries, and explainable grading/recommendation [9].

· **Equity and accessibility**: multilingual and accessibility support **on by default**; availability of bias and safety evaluation reports [2][6].

· **Red-line compliance**: default ban on emotion recognition and profiling-style assessment for minors (aligned with EU AI Act Article 5(1)(f) [11]).

· **Accountability**: clear incident reporting, appeals channels, and lines of responsibility [12].

10.4 Roles and Capacity-Building for Teachers, Parents, and Students

A governance framework with no support from frontline actors remains an abstraction. Teachers are the critical intermediary in the "last mile" of safety and equity governance — they model responsible use, and they are the first to spot risk and the first responders to appeals. This role is already formally institutionalized in international standards: UNESCO's *AI Competency Framework for Teachers*, released in September 2024, is built on principles of protecting teachers' rights, strengthening human agency, and promoting sustainability, and defines the competencies teachers need across 5 dimensions and 15 competencies; the companion *AI Competency Framework for Students* sets out 4 dimensions and 12 competencies spanning a three-tier progression of understanding, applying, and creating, with AI ethics, safety, and meaningful engagement running throughout [6][7]. This provides an operational competency map for treating teacher capacity-building as a first-order investment.

The current state, gaps, and training supply for teacher AI literacy are addressed systematically in Chapter 6 of this book (the dedicated chapter on teacher AI literacy); this chapter emphasizes only its pivotal position in the governance chain: without teachers capable of detecting fabrication, correcting errors, and mediating use, no admission checklist or code of use can truly take hold in the classroom. Parents' informed consent and digital-parenting capacity, together with students' responsible use and self-protection awareness, form the social foundation of governance — particularly under the legal framework requiring guardian consent for children under 14 [15], where parental informed consent is not a formality but a statutory link in the data-protection chain. Data on training coverage and competency assessment appear at [TBD: data/source].

10.5 Conclusion and Implementation Roadmap

10.5.1 Overall Assessment

Looking across the whole book — from literacy, curriculum, and human–AI collaboration through assessment to governance — this Blue Book's core judgment is this: **AI education in primary and secondary schools has evolved from a question of "do we have the resources" to a question of "can we use them equitably, safely, and responsibly."** The opportunities and risks brought by generative and agentic technologies share the same root; what determines the net effect is not the technology itself but the enabling infrastructure — the combined maturity of literacy foundations, curriculum design, human–AI collaboration paradigms, assessment methods, and governance mechanisms.

There is a "capability–governance gap" between technological advancement and governance, and the evidence assembled in this chapter measures that gap clearly. On one side sits the reality documented by UNESCO's survey — fewer than 10% of institutions with formal AI-use policies [8] — alongside the access reality documented by the ITU, with 2.6 billion people still offline [1]. On the other side sits the wave of hard law (the EU AI Act [11], the Council of Europe convention [12], the FTC's COPPA amendments [17]) and soft-law frameworks (the OECD Recommendation [9], UNESCO's guidance and competency frameworks [8][6][7], China's two guidelines and its Minor Protection Regulation [13][14]) enacted in quick succession between 2024 and 2026. Closing this gap — turning frameworks on paper into classroom reality — is the central task of the years ahead.

10.5.2 A Phased Implementation Roadmap

This Blue Book proposes a three-phase roadmap running from 2027 to 2030. What follows is a strategic framework; quantitative targets for each phase should be set by individual jurisdictions against their own local baselines — see [TBD: data/source].

Phase	Time window	Strategic focus	Signature outcomes
Phase 1 · Foundation-building	2027	Scale teacher AI literacy (aligned with the UNESCO teacher framework [6]); implement procurement admission checklists; establish baseline safety codes; enforce red-line compliance (e.g., emotion recognition)	Teacher literacy attainment rate [TBD]; admission-checklist coverage rate [TBD]
Phase 2 · Deepening	2028–2029	Embed AI into curriculum; scale agentic teaching-assistant pilots; transform assessment methods; put accessibility and language equity upstream	Share of pilot schools [TBD]; assessment-reform coverage [TBD]
Phase 3 · Maturity	2030 and beyond	Close the loop on equitable outcomes; normalize governance; sustain evidence-based improvement	Reduction in between-group gaps [TBD]; routine annual assessment mechanism

10.5.3 Recommendations for Each Stakeholder Group

- **For policymakers**: build out K-12 AI governance around "safety baseline + equity upstream + evidence-based assessment," avoiding a situation where regulation lags deployment; write accessibility and language equity into admission as hard constraints; following the EU AI Act and the Council of Europe convention, establish a clear red line banning emotion recognition and profiling for minors [11][12].
- **For education authorities and schools**: build a procurement admission checklist and full-lifecycle governance process; treat teacher capacity-building as the first-order investment, rather than equating device spending with capability-building; make full use of national public-resource platforms to narrow urban–rural gaps at the use layer [5].
- **For vendors and researchers**: hold to "safety by design, accessibility by design" for minors; implement separate consent for collecting minors' biometric data [17]; publish bias and safety

evaluation results; participate in localizing educational data and building out low-resource-language capability.

· **For teachers, parents, and students**: take on, respectively, the responsibilities of modeling and mediating use, giving informed and cooperative consent, and using AI responsibly; parental guardian consent is a statutory link in data protection, not a formality [15] — together these form the social foundation of governance.

10.5.4 Research Outlook and Cross-References Within This Institute

This Blue Book has focused on the capability–governance dimension of the enabling infrastructure; several questions warrant continued tracking: the long-term learning effects and dependency risks of agentic teaching assistants, causal identification of equitable outcomes, and comparable cross-national governance metrics (such as the comparability of "red-line" checklists across jurisdictions). On the broader question of devices and interaction forms, this institute has produced related studies worth cross-referencing: on the application and governance of emerging device-side interaction forms such as AI glasses in education — where continuous collection of minors' biometric and contextual data makes an especially sharp test case for this chapter's red lines on data minimization and banned profiling — see this institute's *2026 Blue Book on the AI Glasses Education Industry*; on the deployment maturity and red lines for educational robotics and embodied AI, see this institute's *Global Blue Book on Educational Robotics Development 2026*. Together, these three studies sketch the full landscape of the enabling infrastructure for K-12 AI education toward 2026 and beyond, characterized by increasing agentic capability, multimodality, and device-side computing.

> **Closing thought.** Technology will keep evolving, but education's fundamental promise does not change: every child deserves to grow up safely, equitably, and with dignity. Whether artificial intelligence becomes the lever that delivers on this promise, rather than a new engine of division, depends on whether we are willing today to write equity into design, establish safety as a baseline, and put governance into practice. With 2.6 billion people worldwide still unable to get online [1], and fewer than one in ten educational institutions yet to adopt a formal AI policy [8], this is at once an unfinished checklist and a clear agenda for action. It is both where this Blue Book arrives and where our shared work begins.

Chapter References

[1] Facts and Figures 2024: Internet Use (global internet use and the digital divide) · International Telecommunication Union (ITU) · 2024 · https://www.itu.int/en/mediacentre/Pages/PR-2024-11-27-facts-and-figures.aspx (see also https://www.itu.int/itu-d/reports/statistics/facts-figures-2024/)

[2] Digital divide in education / PISA 2022 on the socioeconomic gap in digital-resource access · Organisation for Economic Co-operation and Development (OECD) · 2023–2024 · https://www.oecd.org/en/topics/sub-issues/digital-divide-in-education.html

[3] Digital Education Leading the Future — A Review of Positive Progress in China's Education Digitalization [数字教育 引领未来——我国教育数字化工作取得积极成效综述]; and National Survey Report on the Application of Educational Informatization in Primary and Secondary Schools [全国中小学教育信息化应用情况调查报告] (urban–rural coverage rates, hardware/training gaps) · Ministry of Education of the People's Republic of China / China Institute for Educational Financial Research, Peking University · 2024 · http://www.moe.gov.cn/jyb_xwfb/s5147/202401/t20240130_1113352.html; https://ciefr.pku.edu.cn/cbw/kyjb/c6f7200a798748fab33ee2a3ff637ef5.htm

[4] PISA 2022 Results (Volume II): digital-device distraction (65%/30%), phone-related anxiety (45%), the distraction–achievement gap (approximately 15 points), and the positive effect of moderate use · Organisation for Economic Co-operation and Development (OECD) · 2023 · https://www.oecd.org/en/publications/pisa-2022-results-volume-ii_a97db61c-en.html

[5] National Smart Education Platform for Primary and Secondary Schools: cumulative registered users surpassing 100 million, with more than 36.7 billion page views (as of end of 2023); total page views reaching 40.54 billion, with roughly 88,000 resources (as of 2024-05-15) · Ministry of Education of the People's Republic of China · 2024 · http://www.moe.gov.cn/jyb_xwfb/s5147/202405/t20240521_1131711.html; http://www.moe.gov.cn/jyb_xwfb/xw_zt/moe_357/2024/2024_zt02/mtbd/202401/t20240126_1112610.html

[6] AI Competency Framework for Teachers (5 dimensions, 15 competencies) · United Nations Educational, Scientific and Cultural Organization (UNESCO) · 2024 · https://www.unesco.org/en/articles/ai-competency-framework-teachers

[7] AI Competency Framework for Students (4 dimensions, 12 competencies, spanning understanding–applying–creating) · United Nations Educational, Scientific and Cultural Organization (UNESCO) · 2024 · https://www.unesco.org/en/digital-education/g77-competencies (draft framework: https://www.unesco.org/sites/default/files/medias/fichiers/2024/07/UNESCO-Draft-AI-competency-frameworks-for-teachers-and-school-students_0.pdf)

[8] UNESCO survey: Less than 10% of schools and universities have formal guidance on AI (more than 450 institutions surveyed) · United Nations Educational, Scientific and Cultural Organization (UNESCO) · 2023 · https://www.unesco.org/en/articles/unesco-survey-less-10-schools-and-universities-have-formal-guidance-ai

[9] Recommendation of the Council on Artificial Intelligence (adopted 2019, revised 2024; five value-based principles; 38 member countries) · Organisation for Economic Co-operation and Development (OECD) · 2019/2024 · https://oecd.ai/en/ai-principles; https://www.oecd.org/en/topics/ai-principles.html

[10] Guidance for Generative AI in Education and Research (by Fengchun Miao & Wayne Holmes; recommends age-13 minimum, data-privacy protection; addressed to 193 member states) · United Nations Educational, Scientific and Cultural Organization (UNESCO) · 2023 · https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

[11] EU AI Act (European Union Artificial Intelligence Act): education classified as high-risk (Annex III); Article 5(1)(f) bans emotion recognition in educational institutions; Article 27 requires a Fundamental Rights Impact Assessment · European Union · 2024 · https://artificialintelligenceact.eu/annex/3/; https://fpf.org/blog/red-lines-under-eu-ai-act-unpacking-the-prohibition-of-emotion-recognition-in-the-workplace-and-education-institutions/

[12] Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (first binding international AI treaty, signed in Vilnius on 2024-09-05) · Council of Europe · 2024 · https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

[13] Guidelines for General AI Education in Primary and Secondary Schools (2025 Edition) [中小学人工智能通识教育指南（2025 年版）] and Guidelines for Generative AI Use in Primary and Secondary Schools (2025 Edition) [中小学生成式人工智能使用指南（2025 年版）] · Teaching Guidance Committee for Basic Education, Ministry of Education · 2025-05-13 · https://www.cernet.edu.cn/xxh/focus/zc/202505/t20250513_2667990.shtml

[14] Regulation on the Protection of Minors in Cyberspace [未成年人网络保护条例] (State Council Order No. 766; adopted at a State Council executive meeting on 2023-09-20; effective 2024-01-01) · State Council of the People's Republic of China · 2023/2024 · https://www.gov.cn/zhengce/content/202310/content_6911288.htm

[15] Personal Information Protection Law of the People's Republic of China [中华人民共和国个人信息保护法] and Law of the People's Republic of China on the Protection of Minors (as amended 2024) [中华人民共和国未成年人保护法（2024 年修正）]: processing personal information of minors under 14 requires guardian consent, and consent to non-essential processing may not be compelled · Standing Committee of the National People's

Congress · 2021/2024 · https://www.yantian.gov.cn/YTQSF/gkmlpt/content/12/12011/post_12011375.html

[16] Designing for Education with Artificial Intelligence (developer guide, 2024-07); Toolkit for Safe, Ethical, and Equitable AI Integration (school guide, 2024-10) · U.S. Department of Education · 2024 · https://www.ed.gov/about/ed-overview/artificial-intelligence-ai-guidance

[17] FTC finalizes amendments to the COPPA Rule (Children's Online Privacy Protection Rule amendments: biometric data added to personal information; separate opt-in consent required for third-party sharing; effective 2025-06-23) · U.S. Federal Trade Commission (FTC) · 2025 · https://www.k12dive.com/news/ftc-finalizes-coppa-rule-children-data-privacy/738077/

Appendix A Research Methods and Data Conventions

A.1 Research Methods

This Blue Book was produced by AI-SLI using a "self-built curated knowledge base plus AI-assisted production pipeline" methodology. The team first constructed an analytical framework spanning literacy, curriculum, human–AI collaboration, assessment, and governance, then conducted web-based retrieval and evidence-gathering on each topic in turn. Drafting drew solely on verified primary-source material retrieved in this process, with every specific fact tied to an inline citation and the whole subject to human review for directional accuracy and red-line compliance. The research team directed topic selection, framework design, judgment calls, and final conclusions; the AI component carried out intensive retrieval, drafting, tabulation, and citation-management work.

A.2 Data Sources

Data draw primarily on primary sources, including: original policy texts and curriculum standards from national education authorities (such as relevant documents from China's Ministry of Education), frameworks and surveys from international organizations (UNESCO, the OECD, the European Commission), authoritative third-party surveys (such as those from RAND and the Pew Research Center), and publicly available material from vendors and institutions. Each chapter's closing "Sources Cited in This Chapter" section lists, item by item, the genuine sources actually cited in that chapter (title, institution, year, URL).

A.3 Principles of Caution in Data Conventions

- **No cross-survey aggregation**: Because surveys differ in sampling frame, grade-level scope, and definition of "use," their percentages cannot be directly compared or summed across surveys; each citation retains its original convention and provenance.
- **Forecasts kept distinct from actuals**: Any institutional forecast is labeled "forecast" together with its publication date, and is presented separately from actual statistical measurements.
- **Placeholders, never fabrication**: Any specific figure, policy document number, product listing, or citation that could not be verified through retrieval is marked with "[TBD: ...]" — under no circumstances is a value invented. This document is a research edition; all [TBD] items await multi-source verification before the final edition is issued.

A.4 Limitations and Next Steps

This edition retains [TBD] placeholders for certain granular data points — for example, country-by-country instructional hours, specific regulatory clause numbers, and the original-language naming of individual competencies. These will be filled in with verified sources through the fact-checking pipeline in a subsequent final edition.

Appendix B Glossary

Term	Definition
AI literacy	The composite ability to understand, evaluate, and responsibly use artificial intelligence, and to exercise judgment about its social impact.
Generative AI (GenAI)	AI technology capable of generating text, images, speech, code, and other content, epitomized by large language models.
Agent	An AI system capable of understanding goals, decomposing tasks, invoking tools, and executing multi-step tasks continuously.
Multimodal	The capacity to integrate text, images, speech, video, code, and other information forms within a single unified model.
On-device	Deploying inference capability on a personal device or local terminal, with low latency, strong privacy, and offline availability.
AI-native curriculum	A curriculum form reorganized around AI-literacy objectives as its axis, the human–AI collaboration process as its core instructional object, and education-vertical models/agents as its default infrastructure.
Enabling infrastructure	The full complement of capability, curriculum, assessment, governance, and infrastructure elements needed to support AI education in primary and secondary schools.
Prompt engineering	The methods and literacy involved in structuring natural-language interaction with artificial intelligence to obtain reliable output.
Process-oriented assessment	An assessment approach oriented toward the learning process itself (response trajectories, revisions, collaboration, etc.) rather than the final product alone.
Human–AI co-learning	An instructional paradigm in which students and artificial intelligence jointly complete learning tasks as collaborating parties.
Digital divide	Inequality in educational opportunity and outcomes arising from differences in devices, connectivity, computing power, and literacy.
DigComp	The European Union's Digital Competence Framework.
PISA	The OECD's Programme for International Student Assessment.

Appendix C Citation Conventions

This Blue Book follows an "end-of-chapter sourcing" convention: each chapter's closing "Sources Cited in This Chapter" section lists, item by item, the genuine sources actually cited in that chapter (title, institution, year, URL), all of which were actively retrieved and cross-checked as primary sites during production. To avoid duplication, the book does not maintain a separate consolidated bibliography. Wherever "[TBD: ...]" appears in the body text, it marks a genuine source or data point still to be supplied, to be completed in the final edition.

AI-SLI

2026 · AI-Assisted Research Edition