Policy Brief

Series Information:
EPINOVA Policy Brief Series on AI Governance, Strategic Competition, Geopolitics, and Emerging Security

Recommended Citation:
Wu, S. (2026). *AI in China's Higher Education: From National Strategy to Educational Value* (Policy Brief No. EPINOVA–2026–PB–78). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.67037/epinova.pb.2026.078

Disclaimer:
This publication is based on publicly available information, official documents, relevant research, and the author's independent analysis. It does not represent the official views of any government, institution, or organization discussed. The findings and recommendations are intended solely for independent research and policy discussion and do not constitute legal, regulatory, or professional advice.

AI in China's Higher Education:

From National Strategy to Educational Value

Faculty Development, Disciplinary Implementation, and the Evidence Needed to Assess Learning

Author: Shaoyuan Wu

Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC

Date: November 16, 2026

Key Judgments

- **The policy framework is no longer confined to adoption.** China's 2025–2026 education directives link AI literacy, professional development, disciplinary teaching, research infrastructure, educational assessment, and safety. The central implementation question is how these commitments become verifiable practice.
- **Faculty development is a transmission mechanism, not proof of learning impact.** A structured training program can diffuse expectations and expose educators to new methods; course participation and certification do not establish what teachers subsequently implement or what students retain.
- **Educational integration is discipline-dependent, including in safety-critical settings.** The selected cases span language and humanities, social sciences, economics, archaeology, laboratories, and physical education. They require different standards of validity, human oversight, and, where bodily harm is plausible, controls independent of model output.
- **The training cases frame teachers as designers and accountable validators.** Effective integration requires pedagogical selection, source verification, review of model outputs, protection of student autonomy, and responsibility for consequential decisions rather than merely prompt proficiency.
- **AI-enabled output is not equivalent to educational value.** Efficiency gains, interaction counts, and polished assignments differ from durable understanding, independent reasoning, transfer, equity, and safety; verification work and displaced student practice may offset apparent productivity gains.

Executive Summary

China's AI-in-higher-education agenda has entered a phase in which national objectives, faculty development, and discipline-specific teaching demonstrations are increasingly linked in policy design and faculty-development initiatives. The 2025 education-digitalization framework and the April 2026 AI + Education Action Plan do not merely encourage digital tools: they call for more integrated curricula, professional capacity, research platforms, evaluation, and safeguards. A nationally organized faculty-training initiative supplies a concrete, but selective, view of how these ambitions are communicated to university educators.

GLOBAL AI GOVERNANCE RESEARCH CENTER
EPINOVA

Policy Brief

This brief examines the official policy architecture, public notices about the training program, and sixteen course-session summaries in the author's research archive. The sessions range deliberately across writing and social-science courses, archaeology, intelligent laboratories, and physical education, among other contexts. They illustrate a shift in the design space of teaching: from isolated generative-AI tasks toward linked materials, disciplinary data, human–AI workflows, simulation, and feedback. They do not constitute a representative survey of Chinese universities or an independent evaluation of student outcomes.

The main policy challenge is therefore not a missing national commitment to evaluation. The 2026 plan already calls for scientifically assessing educational impact. The challenge is to make this requirement operational at the level of courses, institutions, technology providers, and the students whose learning is at stake. This brief proposes a layered evidence standard, discipline-specific assurance, performance-based teacher development, vendor accountability, public-interest safeguards, and risk-proportionate decision criteria for scaling, revising, or suspending AI-enabled interventions.”

Why This Matters

Higher education develops knowledge, certifies competence, and signals expertise to society. Generative AI can expand access to feedback and complex learning tools, but it also makes assisted performance harder to distinguish from independent mastery. Improved efficiency may conceal errors, shift verification burdens to teachers, or weaken opportunities for independent learning.

AI adoption therefore does not necessarily translate into educational value. Rewarding visible deployment risks prioritizing scale over learning, while overly restrictive evaluation may discourage innovation. Policymakers must assess actual learning outcomes, costs, and risks rather than presume success or failure.

The implications extend beyond universities. Governments need evidence to allocate resources; educators need reliable assessments; industry needs clear standards; and students need equitable access, privacy, and accountable decisions. **Effective policy must balance innovation with verifiable educational outcomes and safeguards.**

Analytical Scope and Evidence Rules

This brief examines how China's AI education policies translate into faculty training and teaching practices. It analyzes policy documents and selected training examples without estimating national adoption rates or causal educational effects.

Sources include official policies (Central Committee of the Communist Party of China & State Council of the People's Republic of China, 2025; Ministry of Education et al., 2025, 2026), university training notices (Guilin University of Electronic Technology, 2026; Henan Normal University, 2026), sixteen training-session summaries (Wu, 2026), and international guidance (UNESCO, 2023, 2024). Faculty adoption studies provide additional context but not evidence of learning outcomes (Gao & Yang, 2026; Luo & Day, 2026).

Policy intentions do not establish implementation; training completion does not prove teaching competence; and AI-assisted performance does not demonstrate independent learning. Proposed validation scenarios are illustrative, not documented failures. Findings identify implementation patterns and evidence requirements, not nationally representative outcomes.

Policy Brief

1. From National Strategy to Institutional Implementation

1.1 Policy Objectives and Their Administrative Reach

The 2025 *Outline for Building an Education Powerhouse* places educational quality, digital modernization, talent development, and strong higher education within a long-term state strategy. The nine-department *Opinions on Accelerating Education Digitalization* further link high-order thinking, judgment, practical ability, digital resources, data governance, and AI-supported education (Central Committee of the Communist Party of China & State Council of the People's Republic of China, 2025; Ministry of Education et al., 2025).

The April 2026 *AI + Education Action Plan* gives this agenda a more operational architecture. Among its provisions are AI foundational learning and interdisciplinary curricula in universities; intelligent teaching tools and research agents; smart laboratories and future learning spaces; infrastructure for computing, data, and models; mechanisms for evaluating educational effects; and safety controls. Its 2030 objective is an intended future state, not an observed achievement in 2026 (Ministry of Education et al., 2026).

These measures combine several policy logics. AI is treated as a subject of learning, a means of teaching, an input to research, an infrastructure procurement problem, and an object of governance. A higher-education evaluation that measures only model access or prompt use would therefore miss much of what policy is trying to accomplish. Conversely, a centrally specified platform cannot by itself establish whether a course's disciplinary standards have been met.

1.2 Faculty Training as a Policy Transmission Channel

Public university notices describe a 2026 online training pathway involving a competency assessment, course recommendations, three training streams, examinations, and certificates. The streams are Future Teachers; one of four discipline-specific options—humanities/social sciences, science and engineering, agriculture, or medicine; and experimental teaching. Some notices specify a minimum of sixteen courses across the streams. A certificate in the author's archive records sixteen training hours; these are different administrative measures and should not be conflated (Guilin University of Electronic Technology, 2026; Guilin Medical University, 2026; Wu, 2026).

The assessment materials identify seven dimensions of future-teacher development: ethics and educational values; teaching innovation and design; student development; integration of research and knowledge creation; digital literacy and human–AI collaboration; social responsibility and global perspective; and resilience and lifelong professional learning (Wu, 2026). This is a curriculum and self-assessment architecture, not an externally validated score for the national teacher population.

For policy analysis, training is an intermediary institution. It can standardize vocabulary, expose faculty to possible workflows, and create channels for disseminating examples. What happens afterwards depends on institutional incentives, disciplinary resources, available time, pedagogical judgment, and student needs. The relevant next-stage evidence is a documented course redesign and independently reviewed implementation rather than a certificate count alone. A completion-oriented administrative system could also reward credentials or module passage more visibly than sustained classroom change, creating a risk of surface compliance. This is a proposed incentive mechanism, not an empirical finding about the training platform or its participants; testing it would require subsequent classroom artifacts, participant interviews, and

follow-up assessment. The adoption literature reinforces the need to examine institutional support and teacher autonomy alongside technical proficiency (Gao & Yang, 2026; Luo & Day, 2026).

2. What the Cross-Disciplinary Cases Reveal

2.1 Beyond Standalone Chatbots

The sixteen session summaries include AI-supported writing feedback, course knowledge graphs, teaching assistants, intelligent MOOCs, microeconomic decision exercises, research-agent workflows, virtual laboratories, and multi-agent experimental projects. Rather than a single technology, the material describes ways to connect course content, resources, student activity, teacher review, and assessment into repeatable instructional processes (Wu, 2026).

This suggests an important change in the unit of evaluation. Testing whether a model can answer a question is different from testing whether a semester-long course workflow improves learning. A workflow can improve feedback speed while increasing teachers' verification burden or exposing student data to third parties. Evaluation must therefore identify the intervention's boundaries and counterfactual: a prompt, a feedback loop, a course unit, a shared platform, or an institution-wide system.

2.2 Different Disciplines, Different Validity Standards

The cases do not support a one-size-fits-all model of AI-enhanced education. In archaeology and humanities research, the critical test is provenance: whether the cited artifact, image, record, or interpretation can be independently traced. In economics and social-science teaching, the test concerns assumptions, counterexamples, identification, and the distinction between simulation and observed social behavior. Writing and language courses add audience, intercultural nuance, independent revision, and the student's ability to explain decisions (Wu, 2026).

Laboratory and physical-education examples require another kind of assurance. An apparently plausible simulation can misrepresent real equipment; a posture or movement-recognition tool may fail differently across individuals. For live experiments, physical equipment, exercise prescriptions, and other safety-relevant decisions, AI outputs should not directly actuate hazardous processes or prescribe consequential action without a qualified human's authorization. Where risk assessment identifies severe or hard-to-reverse consequences, independent physical interlocks, readily available emergency stops, and two-person approval for specified hazardous steps are proportionate safeguards; ordinary low-risk simulations need not inherit the same controls. A safety-critical failure should suspend the affected AI-controlled function, not automatically every course activity, until investigation and reauthorization. These are proposed protections consistent with risk-based AI management, not observed failures in the sessions (National Institute of Standards and Technology [NIST], 2023; Wu, 2026). The 'Safety Culture' session contributes a risk-governance perspective; it is not evidence that AI improved laboratory safety.

These distinctions imply that shared rules for privacy, accountability, source disclosure, and appeal should coexist with discipline-specific evaluation profiles. The national objective can be common even when the meaning of a valid educational result differs substantially between a humanities archive, a chemistry laboratory, a field-sport exercise, and a social-science simulation.

2.3 Two Illustrative Micro-Level Validation Tests

- **Test A — Microeconomics: Plausibility Without Model Validity.** A session on microeconomics uses AI-assisted preparation and examples involving pricing and labor-market reasoning (Wu, 2026). A targeted stress test would ask students to evaluate a polished AI explanation of minimum-wage effects that silently assumes a perfectly competitive labor market even though the problem specifies employer market power. The key evidence would be whether students identify the missing assumption, supply a counterexample, and defend a corrected argument **without** AI. Teachers must review economic assumptions and the scoring rubric; the output's fluency and time saved do not establish learning. **This is a proposed test, not a recorded classroom incident.**
- **Test B — AI-Assisted Archaeology: Candidate Detection Without Ground Truth.** The archaeology session discusses computational identification of possible sites and related interpretation (Wu, 2026). A proposed audit would compare model-generated candidates with available independent archaeological records, expert review, and, where lawful and practicable, field confirmation. The count of candidate locations is a workflow output, not a confirmed count of archaeological discoveries. The verification workload, false-positive rate, and uncertainty labels must be recorded before claims of improved identification accuracy. **The archive does not establish that the session's model made a specific false identification.**

2.4 The Teacher as Designer, Verifier, and Decision Owner

Across the teaching-assistant, backward-design, language, and experimental sessions, the teacher's proposed role shifts from presenting completed material to specifying learning goals, choosing what AI may do, designing student interactions, judging evidence, and accepting responsibility for the grade or intervention (Wu, 2026). This is not proof that every institution has achieved that role transformation; it is a recurring model in the materials.

UNESCO's *AI Competency Framework for Teachers* connects AI literacy with a human-centred mindset, ethics, AI foundations, pedagogy, and professional learning (UNESCO, 2024). Its comparison with the seven-domain Chinese training assessment reveals overlapping concerns but different categories and purposes. Neither should be reduced to an interchangeable checklist. The practical issue is whether training assessments ask teachers to demonstrate judgment: detecting plausible misinformation, refusing an unsuitable tool, designing an accessible alternative, and justifying how a student will show independent mastery.

3. From AI Activity to Educational Value

3.1 Four Evidence Layers and Cross-Cutting Safeguards

AI education evaluation should distinguish four evidence layers: **access and participation, implementation fidelity, assisted task performance, and durable independent capability**. While participation and efficiency indicate adoption, they do not establish learning gains. Independent assessments, delayed retention, and knowledge transfer provide stronger evidence of educational effectiveness.

Educational value should be assessed through durable learning gains, independent capability, and equitable access, alongside verification costs and potential harms. Privacy, accessibility, and safety should be evaluated across all four layers. Potential negative externalities, including reduced independent practice, increased teacher verification burdens, and unequal access to feedback, must also be considered. These remain risks to be tested, not demonstrated outcomes of the sixteen training sessions (Ministry of Education et al., 2026; UNESCO, 2023).

Table 1. Minimum Evidence for Different Educational Claims

Claim or Decision	Suitable Evidence	What It Does Not Establish
Access and Uptake	Eligible students, participation denominator, costs, accessibility, and recorded tool usage	Learning gain or fairness
Implementation Fidelity	Course protocol, model/version, actual teacher review, student task sequence, and deviations	Independent mastery
Assisted Performance	Rubric-scored work with AI permissions recorded and human moderation	Performance without AI
Durable Capability	Baseline, comparable independent task, delayed assessment, unfamiliar transfer task	Universal effects outside tested settings
Safety and Equity	Subgroup access, privacy incidents, error appeals, accessibility audits, adverse events	Absence of unobserved harm
Net Educational Value	Changes in independent-task ability, displaced practice time, teacher verification hours, and complaint burden	A universal cost-benefit result

Source. Author's proposed assessment framework, informed by national policy and international guidance (Ministry of Education et al., 2026; UNESCO, 2023). The categories are not reported national requirements.

3.2 The Risk of Misclassification

There are at least three ways to overstate success. First, exposure can be mistaken for expertise: a teacher completing a course does not demonstrate competent application. Second, output quality can be mistaken for learning: an AI-assisted essay may be excellent even if the student cannot reconstruct its argument. Third, platform scale can be mistaken for equity or value: increased access may coexist with gaps in connectivity, disability accommodation, disciplinary data, or human feedback.

A stronger design would record the pre-intervention baseline, the actual AI permissions, comparison conditions when feasible, scoring methods, teacher labor, student time spent independently searching and reasoning, and follow-up outcomes. Randomization or staggered rollout may be viable for some course pilots, but not for all. When causal identification is not possible, credible reporting should mark the study as descriptive and specify its limitations rather than convert association into effectiveness. The 2026 policy framework already creates room for scientific assessment; the policy problem is how consistently such distinctions can be enacted and publicly documented (Ministry of Education et al., 2026).

4. Policy Implications Across Six Levels

4.1 National Policy and Intergovernmental Coordination

Decision: Whether to move from promoting exemplary AI cases toward certifying evidence quality as a condition for expansion. The central government can define a minimal reporting schema aligned with existing evaluation ambitions: learning objective, unit of intervention, baseline, tool/version, accessibility, human review, measured outcome, and failure reporting. Provinces can coordinate pilots while allowing institutions to adapt disciplinary tests. Public funding should distinguish infrastructure acquisition from evidence of educational return and examine regional differences in support capacity (Ministry of Education et al., 2025, 2026).

Decision Trigger: Nationally promoted exemplar programs should be required to publish an auditable evaluation synopsis before being treated as transferable models. A missing baseline or missing independent learning measure should lead to a 'promising practice—not yet validated' classification, not automatic rejection or automatic scale-up. This is a proposed rule, not a description of an existing nationwide certification system.

4.2 Universities, Teaching, and Academic Assessment

Decision: Which parts of teaching should be assisted, redesigned, or left deliberately independent. University academic affairs offices should require course-level statements of permissible AI use, discipline-specific verification, alternate routes for students who cannot use the tool, and a defensible assessment that separates supported output from independently retained capability. Faculty development centers can evaluate teacher performance through implemented artifacts and follow-up reflection rather than module completion alone (Wu, 2026; UNESCO, 2024).

Decision Trigger: Before a course claims improved learning, it should provide an aligned rubric, a baseline, a teacher-review protocol, and at least one independent or delayed assessment. Failed cases should be retained as learning evidence. For high-stakes grading, students should have an accessible route to human review and appeal.

4.3 Technical Infrastructure, Models, and Data Governance

Decision: Whether a technically functional system is a dependable educational intervention. Institutions should document model versions, retrieval sources, data flows, access controls, known failure modes, and the procedures for correcting hallucinated citations, unsafe experimental advice, or biased feedback. Procurement should prefer interoperable exports and auditable interfaces where feasible. Contracts should require advance notice of material model, alignment, safety-filter, retrieval-index, data-processing, or scoring changes, with version pinning or rollback where technically available. Because some hosted models change without an exposed version identifier, institutions should also retain a frozen regression set and monitor sentinel outputs for drift; any significant change triggers proportional revalidation before consequential uses. Such controls reduce but do not eliminate the risk of silent updates. Data minimization and transparent retention rules are especially important where platforms infer student profiles (Ministry of Education et al., 2026; NIST, 2023; UNESCO, 2023).

Decision Trigger: Material changes in the model, alignment behavior, instructional or retrieval data, assessment prompt, or data-processing vendor should initiate a documented regression test. If the supplier cannot disclose versions or provide a safe rollback, the affected high-stakes function should remain in a human-only or validated fallback mode until testing is complete. A vendor benchmark alone is insufficient to establish educational validity.

4.4 AI Industry, Education Technology, and Employers

Decision: Which outcomes suppliers and labor-market partners should be asked to demonstrate. Education-technology firms can be held to procurement declarations about provenance, privacy, bias, accessibility, uptime, human override, and outcomes appropriate to the purchased use case. Employers and professional bodies can help identify new skill demands and supply authentic problems for learning, but should not define academic validity on their own. China's broader *AI+* initiative connects AI with industry and social transformation, making this boundary between labor-market relevance and educational autonomy particularly consequential (State Council of the People's Republic of China, 2025).

Decision Trigger: Contract renewal or scaled procurement should consider independent usability, learning, fairness, and security evidence rather than merely license counts or product demonstrations. Partnerships should disclose commercial interests and preserve university control of grading, credentialing, and research conclusions.

4.5 Society, Learners, and Equal Opportunity

Decision: How to protect the public value of credentials and the learner's ability to exercise independent judgment. Students in less-resourced institutions, learners with disabilities, and students with different language or digital-literacy backgrounds may face different gains and risks. Institutions should therefore measure not only who received AI access but who received usable support and who bears the burden of error correction. This includes *task-formulation and verification literacy*: the capacity to pose a well-bounded question, challenge AI-generated answers, locate trustworthy disciplinary evidence, and revise work. Differences in prior preparation, language support, and mentoring may make nominally equal access unequal in practice; this possibility requires subgroup data rather than assumptions about disadvantaged students. A human-centred model requires meaningful notice, proportional data collection, accessible alternatives, and appeal channels for AI-influenced academic decisions (UNESCO, 2023, 2024).

Decision Trigger: Persistent subgroup gaps, unremedied accessibility failures, or repeated consequential misgrading should trigger an immediate audit and targeted remedy. Suspend the affected use case when material harm is ongoing, credible, and cannot be contained through prompt safeguards; do not presume any disparity alone proves discriminatory model causation. Public reporting should preserve privacy and avoid reducing personal development to behavioral surveillance.

4.6 Measurement and Statistical Governance

Decision: What can legitimately be compared across campuses, courses, and time. A common statistical reporting protocol should specify denominators, missing data, exposure intensity, changes in assessment instruments, nonrandom selection, confidence intervals where justified,

and the distinction between descriptive, associational, and causal results. Institutions should report negative and mixed findings and avoid ranking universities on raw AI usage or unverifiable 'AI-enhanced' course counts. The adoption studies of Gao and Yang (2026) and Luo and Day (2026) illustrate that carefully designed surveys can establish correlates of faculty behavior, but these outcomes should not be relabeled as student learning effects.

Decision Trigger: Any cross-campus comparison should identify whether it compares like-for-like units and whether its instruments are sufficiently invariant. Where that standard is not met, publish a non-ranking descriptive account instead of a misleading performance league table.

Limitations

The sixteen training-session summaries constitute a purposive archive, not a representative sample of Chinese universities or educators. Policy documents and training notices establish stated objectives and procedures, not national adoption or completion rates (Guilin University of Electronic Technology, 2026; Wu, 2026).

The original sessions have not been fully re-audited, and subsequent classroom implementation, platform records, and incident data remain unverified. Illustrative scenarios and proposed safeguards are analytical recommendations, not documented failures. Training certification does not establish teaching effectiveness.

AI usage, productivity gains, and assisted outputs do not establish independent learning, equity, or safety. This brief estimates no causal or nationwide effects; its evidence gaps concern only the materials examined. UNESCO and NIST provide normative references, while external adoption studies do not validate educational outcomes (Ministry of Education et al., 2026; NIST, 2023; UNESCO, 2024).

Conclusion

China's higher-education AI strategy is notable for connecting national objectives, teacher development, discipline-specific teaching, research modernization, and safety within one policy architecture. The sixteen training observations provide an informative view of the kinds of practices being promoted, but not a verdict on nationwide adoption or student achievement. What is observable is a widening set of educational applications and a growing demand for teacher judgment, disciplinary verification, and institutional coordination.

The decisive policy question is whether expansion can be accompanied by credible evidence that students learn more, retain what they learn, transfer it to unfamiliar problems, and exercise judgment even when AI is available. Such evidence cannot be replaced by certification counts, platform interactions, model benchmarks, or impressive demonstrations. Nor should a demand for evidence become a blanket prohibition on innovation. A disciplined, risk-proportionate regime can protect experimentation by distinguishing promising pilots from proven outcomes, making failure visible enough to correct, and immediately isolating genuinely safety-critical failures without suppressing low-risk educational exploration.

The next phase should be judged less by the number of institutions saying they use AI than by the clarity of the educational claims they can substantiate, and by the choices those claims enable governments, universities, industry, and society to make.

References

Central Committee of the Communist Party of China & State Council of the People's Republic of China. (2025, January 19). *Outline for building an education powerhouse (2024–2035)* [教育强国建设规划纲要（2024—2035年）]. https://www.moe.gov.cn/jyb_xxgk/moe_1777/moe_1778/202501/t20250119_1176193.html

Gao, J., & Yang, W. (2026). Faculty adoption of AI-assisted teaching tools in Chinese higher education: An integrated UTAUT2–TAM framework at Shanghai University. *Frontiers in Artificial Intelligence, 9*, Article 1847439. https://doi.org/10.3389/frai.2026.1847439

Guilin Medical University. (2026, November 3). *Notice on the 2026 AI-empowered talent development in higher education faculty training series* [关于开展2026年度“人工智能赋能高等教育人才培养”系列师资培训的通知]. https://www.glmu.edu.cn/jsfzzx/info/1034/2730.htm

Guilin University of Electronic Technology. (2026, November 2). *Notice on the 2026 AI-empowered talent development in higher education faculty training series* [关于开展2026年度“人工智能赋能高等教育人才培养”系列师资培训的通知]. https://www.guet.edu.cn/utsc/2026/1102/c5309a161862/page.htm

Henan Normal University. (2026, October 29). *Notice on participating in the 2026 national faculty training series* [关于组织参加教育部“人工智能赋能高等教育人才培养”系列师资培训的通知]. https://www.htu.edu.cn/jszx/2026/1028/c23439a409871/page.htm

Luo, Y., & Day, M. J. (2026). Determinants of lecturer readiness to adopt generative AI in higher education: Survey evidence from UTAUT and self-determination theory. *Education and Information Technologies, 31*, 3399–3430. https://doi.org/10.1007/s10639-026-13931-3

Ministry of Education of the People's Republic of China, Cyberspace Administration of China, National Development and Reform Commission, Ministry of Science and Technology, Ministry of Industry and Information Technology, Ministry of Public Security, Ministry of Finance, Ministry of Human Resources and Social Security, & National Data Administration. (2025, April 11). *Opinions on accelerating education digitalization* [关于加快推进教育数字化的意见]. https://www.moe.gov.cn/srcsite/A01/s7048/202504/t20250416_1187476.html

Ministry of Education of the People's Republic of China, National Development and Reform Commission, Ministry of Industry and Information Technology, Ministry of Science and Technology, & National Data Administration. (2026, April 2). *AI + education action plan* [“人工智能+教育”行动计划]. https://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

State Council of the People's Republic of China. (2025, August 21). *Opinions on deeply implementing the “AI+” initiative* [国务院关于深入实施“人工智能+”行动的意见]. https://www.cac.gov.cn/2025-08/27/c_1758018277755538.htm

UNESCO. (2023). *Guidance for generative AI in education and research*. https://unesdoc.unesco.org/ark:/48223/pf0000386693

UNESCO. (2024). *AI competency framework for teachers*. https://unesdoc.unesco.org/ark:/48223/pf0000391104

Wu, S. (2026). *AI-empowered talent development in higher education: Sixteen training-session summaries and future-teacher competency assessment materials* [Unpublished research notes]. EPINOVA LLC.