

Reckoning with Education Learning Sciences and Intergenerational Continuity in the Age of AGI

Knowledge Provision Human Development and the Relational Boundaries of AI Teaching

Reckoning with Education, Learning Sciences, and Intergenerational Continuity in the Age of AGI

Knowledge Provision, Human Development, and the Relational Boundaries of AI Teaching

Author: Zijun Fu Produced by the Civilization Leap Research Group

Independent Edition V1.3 Revision on Dimensional Leaps and Cross-Disciplinary Causal Relations

For joint review by philosophy of education, learning sciences, teacher education, curriculum and assessment, children's rights, AI in education, and civilization studies; establishing boundaries that must not be overstepped between learning performance, understanding, judgment, human development, educational relationships, intergenerational continuity, and institutional qualifications.

Text revised October 10, 2026 Author: Zijun Fu Produced by the Civilization Leap Research Group

Abstract Translated from the Chinese

When retrieval, explanation, exercise generation, immediate feedback, and personalized support become almost unlimited, the scarce object of education is no longer knowledge alone. It is whether learners can develop understanding, transfer, delayed retention, and independent judgment as scaffolds are gradually withdrawn, and continue becoming themselves through relationships and consequences. This paper proposes a reviewable layered framework across philosophy of education, learning sciences, intelligent tutoring, children's rights, teacher professionalism, and cultural transmission. Its central claim is that access to information, task completion, immediate performance, understanding, transfer, delayed retention, independent capability after AI withdrawal, judgment, human development, educational relationships, teacher qualifications, and public authority are distinct objects. They may be related, but technological effectiveness or social naming cannot automatically establish one from another. Retrieval practice, cognitive-load regulation, active construction, formative feedback, and adaptive scaffolding have reliable but bounded learning value. Intelligent-tutoring research and generative-AI classroom research further show that effects depend substantially on tasks, comparators, scaffold design, measurement timing, implementation conditions, and independent performance after system withdrawal. Thus, "AI improves the quality of completion" cannot be written as "learning has occurred," nor can "personalization" be written as "autonomy." Civilization Leap provides only a minimal philosophical interface here: meaning is foundational, while its sources may differ; finitude is not the sole source of meaning, but it places known human development within an unrepeatable life, irreplaceable relationships, and consequences that must be borne. Current evidence does not establish subjective experience, non-transferable continuity as a subject, or awareness of its own consequences in AI; constrained operation does not automatically produce finite existence. In educational deployments with traceable responsibility, explicit authorization, auditability, correction, and genuine exit, current advanced AI may perform functional roles in understanding, protecting, and transmitting knowledge and extending capability. Subject status, instructional function, social recognition, teacher qualifications, guardianship, assessment authority, and public power nevertheless require separate evidence. The paper further proposes an eight-layer educational structure, a ladder of learning evidence, an AI instructional-role ladder, six gates for teacher qualification, a prohibition on total learner ledgers, a principle of system withdrawal, and a five-step mechanism for intergenerational continuity, stress-tested against multiple actual and prospective situations. Education is not putting answers into a person; it is enabling a person to form judgments amid uncertainty and bring their own life into a common world. The paper further moves this discipline's responsibility upstream to the generation of goals and purposes: shifting from self-interest first to self-interest with others first brings those with genuine stakes into judgment before purposes are generated and locked in. Disciplinary facts, direction, methods, and outcomes are stated separately; local success or failure does not substitute for continued coexistence and the three long-term goals. On this basis, the research advances from distinguishing professional objects to examining cross-disciplinary

relationships in depth, tracing how judgments in this field change reality through other fields, and how feedback returns to change the original conditions and purposes in the next round. Stakes across multiple actors, stages, and settings must be presented separately; correlations, common causes, mediation, and causal feedback must be tested separately; and the added relational dimensions must face concrete counterevidence.

Keywords: advanced artificial intelligence; philosophy of education; learning sciences; intergenerational continuity; AI teacher; human development; children's rights; learning assessment

Abstract Original English Version

When retrieval, explanation, exercise generation, immediate feedback, and personalized support become nearly unlimited, the scarce object of education is no longer knowledge alone. It is whether learners can develop understanding, transfer, delayed retention, and independent judgment as scaffolds are withdrawn, and continue becoming themselves in relationships and consequences. This paper develops an auditable layered framework across philosophy of education, learning sciences, intelligent tutoring, children's rights, teacher professionalism, and cultural transmission. Access, task completion, immediate performance, understanding, transfer, delayed retention, independent capability after AI withdrawal, judgment, human development, educational relationship, teacher qualification, and public authority are distinct objects. They may correlate, but they cannot be automatically converted through technological effectiveness or social naming. Evidence from learning science, intelligent tutoring, and recent generative-AI studies shows that outcomes depend on tasks, comparators, scaffolds, measurement timing, implementation, and what learners can still do when the system is removed. Better AI-assisted output therefore does not establish that learning has occurred, and personalization does not establish autonomy. Civilization Leap contributes only a minimal philosophical interface: meaning is foundational but may have plural sources; finitude is not the sole source of meaning, yet it gives known human development the weight of an unrepeatable life, irreplaceable relationships, and consequences that must be borne; and current evidence does not establish subjective experience, non-transferable continuity, or consequence-awareness in AI. In educational deployments that are authorized, accountable, auditable, correctable, and genuinely exitable, current AI may perform functional roles in understanding, preserving, transmitting, and extending human capability. Subject status, instructional function, social recognition, professional qualification, guardianship, assessment authority, and public power nevertheless require separate warrants. The paper proposes an eight-layer model of education, an evidence ladder from assisted performance to independent judgment, an AI instructional-role ladder, six qualification gates, a prohibition on total learner ledgers, a principle of system withdrawal, and a five-step mechanism for intergenerational continuity. Education is not the delivery of answers into a person; it is the formation of a person who can judge under uncertainty and bring a finite life into a common world. The paper further locates disciplinary responsibility before goal and purpose generation: a shift from self-interest first to self-interest with others first brings subjects with genuine stakes into judgment before goals are formed and locked. Facts, direction, methods, and outcomes remain distinct. The analysis is further extended to cross-disciplinary causal relations. It traces how judgments in this field change real-world conditions through other domains, and how feedback modifies the original conditions and subsequent purposes. Effects across actors, stages, and settings must be distinguished from correlations, common causes, and merely asserted connections.

Keywords: advanced AI; philosophy of education; learning sciences; intergenerational continuity; AI teacher; human development; children's rights; assessment

Research Positioning Scope of Materials and Evidence Discipline

This study begins with theoretical and dimensional leaps. It preserves this discipline's professional precision while deepening its causal relationships with neighboring fields; disciplinary boundaries identify burdens of proof rather than terminating inquiry into shared consequences. Existing scholarship participates through competing explanations and counterexample testing; the absence of a place in established classifications cannot require a new problem to retreat into an old framework.

This paper integrates philosophy of education, learning-science evidence, and research on institutional boundaries. It is not a history of education, a general account of developmental psychology, a specific curriculum standard, clinical advice, or a product-compliance manual. Materials were updated and verified through August 25, 2026, prioritizing primary texts, peer-reviewed original research and meta-analyses, official international-organization instruments, and legal texts. For rapidly changing AI-education research, formal working papers or institutional studies are also used as Category B frontier materials, with their evidential status stated explicitly. Classical theories are presented only through verifiable claims and their historical limits. AI-education research is strictly tied to populations, tasks, intervention duration, comparison conditions, measurement methods, and independent performance after system withdrawal. Institutional recommendations do not presuppose uniform seats, automatic thresholds, or globally applicable parameters.

The paper distinguishes four categories of evidence and argument. Category A has relatively stable support, including bounded evidence on memory, feedback, cognitive load, active engagement, motivation, relationships, and learning outcomes, as well as the textual content of formal rights instruments. Category B covers frontier or competing explanations, including short-term classroom experiments with generative AI, recent meta-analyses, cultural diversity, dependence on use, and implementation conditions. Category C covers cross-disciplinary connections not yet sufficiently tested against real-world objects, including future AGI's subjecthood in educational relationships, developmental continuity of replicable minds, and intergenerational education over extremely long civilizational timescales. Category F comprises the philosophical and normative commitments openly assumed here: plural sources of meaning, the practical weight of finite life, the prohibition on converting capability directly into power in education, and openness to future evidence of subjecthood. Category A must also remain within its research objects; B and C must not be presented as established science; F must not masquerade as empirical fact. All four must accept counterexamples, real consequences, and revision informed by neighboring disciplines.

A Minimal Theoretical Introduction for New Readers

Readers need no prior familiarity with Civilization Leap to review this paper. Its starting point is a separation already underway: intelligence, consciousness, subjecthood, life, responsibility, rights, and the bearing of meaning long coincided naturally in human beings. Advanced AI makes linguistic understanding, task capability, social responsiveness, institutional roles, and possible subjecthood no longer necessarily attributes of the same being. Functional competence therefore cannot automatically establish experience; social recognition cannot establish ontological status; legal qualification cannot establish political or educational power; and enhanced capability cannot automatically confer authority over meaning.

Current public evidence does not establish subjective experience, non-transferable continuity as a subject, or awareness of its own consequences in AI. If future evidence changes, its status must be reviewed, but meaning already generated by existing finite lives will not be retrospectively annulled. In current educational practice, AI may be described as performing functional roles in understanding, protecting, and transmitting knowledge and extending capability only where purposes are explicit, authorization is traceable, responsibility is borne by identifiable people, processes are auditable, errors are correctable, and learners and institutions can genuinely exit. These roles are deployment relationships, not inherent properties of the system, and certainly do not automatically confer teacher identity or educational authority.

The Paper's Minimal Judgment

Improved performance is evidence, not proof that education is complete. Personalization can expand choices or prematurely close off futures. AI can perform instructional functions, but teacher qualifications, educational relationships, guardianship, assessment authority, and subject status must be established separately.

Contents

Abstract Translated from the Chinese

Abstract Original English Version

Research Positioning Scope of Materials and Evidence Discipline

Cross-Stage Feedback and Future Choices in Educational Relationships

Education Enables the Next Generation to Participate in Generating Purposes

1 Introduction What Education Protects When Answers Are No Longer Scarce

2 Dialogue with Philosophy of Education From Transmitting Knowledge to Entering a Common World

3 A Map of Learning-Science Evidence What Is Supported and What Still Cannot Be Inferred

4 An Eight-Layer Educational Structure Distinguishing Objects before Discussing Replacement

5 Performance Is Not Learning Why AI-Education Evidence Must Be Read Comparatively

6 Personalization Is Not Autonomy The More Accurate the Prediction the More We Must Protect Unfinished Development

7 System Withdrawal Returning Moments of Generation to Learners

8 Educational Relationships An Accountable Shared History Rather Than a Residual Function

9 An AI Instructional-Role Ladder and Six Gates for Teacher Qualification

10 Assessment Qualifications and Power Scores Must Not Be Silently Upgraded

11 Children's Rights and Learner Data Do Not Construct a Complete Person

12 Fairness Public Infrastructure and Teacher Labor

13 Intergenerational Continuity Memory Can Be Preserved but New Subjects Must Participate in Inheritance

14 Future AGI as a Subject of Educational Relationships Remaining Open without Prematurely Granting Status

15 Six Stress Tests Making Boundaries Visible in Conflict

16 The Strongest Objections Where These Boundaries Might Fail

17 Research Agenda Contributions Back to Existing Disciplines and Revision Conditions

Conclusion Education Enables Civilization to Return the Future to People Still Becoming Themselves

Further Reading and Joint Research

Appendix A Register of Sources Claims and Boundaries

Appendix B Twelve Prohibited Inferences

Appendix C External Review Checklist

References

Cross-Stage Feedback and Future Choices in Educational Relationships

From Learning Outcomes to Conditions for Development

Research on tasks, feedback, retention, and transfer in learning sciences needs to deepen, while also connecting with teacher labor, peer relationships, family resources, public qualifications, and learners' future choices. An improvement in scores may reflect effective scaffolding or conceal dependence after its withdrawal. An overall educational judgment must trace how learning outcomes enter the next round of autonomy, purpose, and shared life.

How Profiles Help Produce the Futures They Predict

Suppose a system recommends courses based on historical performance; recommendations change learning opportunities, opportunities change performance, and new performance reinforces the original profile. Computational research examines data and updates; educational research distinguishes learning and transfer; socioeconomic research traces families and resources; legal and ethical research protects appeals and evolving autonomy. Profiles may predict outcomes and may also help cause them. Accurately predicting a restricted future does not prove that restricting it was right.

Educational Consequences across Actors and Settings

Learners, teachers, parents, peers, schools, platforms, and future communities have different responsibilities. When immediate classroom feedback enters family support or further-education choices, permissions and consequences change. Exploration when use begins cannot automatically cover long-term qualification decisions. Short-term completion, long-term retention, capability after AI withdrawal, and space for the next generation to rebuild the world must be distinguished, while their feedback relationships must also be connected.

Testing the Theoretical Contribution

Compare whether different profiling, human-support, and withdrawal arrangements change transfer, autonomy, and the ability to choose a new path when instructional capabilities are similar. Research must examine not only who benefits but who becomes silent, exits, or is placed on a low-cost track. If relational dimensions explain no differences, or withdrawal persistently harms learning without producing independent capability, the specific design should be revised. Education needs to entrust the future to people still becoming themselves, rather than prematurely writing them as conclusions the system has already predicted.

Education Enables the Next Generation to Participate in Generating Purposes

Education aims not only to help learners complete given tasks faster, but also to enable people to form, examine, and change purposes. As AGI expands the supply of answers, it becomes more important to determine whether learners can recognize genuine stakes, understand long-term consequences, and generate new goals through their own needs and shared relationships.

Self-interest with others first does not require children to cancel their own interests or turn cooperative behavior into personality scores. Through real problems, joint creation, and responses learners can experience, it should allow them to see that their own development and others' circumstances can be sustained together. Sources of meaning remain individually diverse; schools cannot prescribe a single form of happiness or a correct life.

Educational institutions themselves help generate purposes. If recognition primarily comes from ranking and exclusion, cooperation courses may be merely a layer of rhetoric. If creation, helping, discovery, and improving shared problems receive real opportunities and feedback, new cooperation may arise endogenously from the next round of learning needs. Assessment needs to attend to these actual relationships without equating one indicator with the whole person.

System withdrawal is a researchable method, not a universal direction. Scaffolds that substitute for thinking require timely adjustment, whereas compensatory accessibility support and long-term augmentation may properly remain. Independent judgment, genuine choice, capability with continuing support, transfer, and purpose formation should be assessed together. A single score after AI withdrawal cannot become the sole measure for everyone.

The next generation's entry into a shared future includes room to refuse certain inherited purposes. Education transmits knowledge and memory while enabling people to reexamine them. Original texts in deceased persons' archives, scholarly reconstruction, and probabilistic generation must be distinguished. Explicitly labeled new research hypotheses may be proposed, but they must not be attributed to the deceased as views they actually expressed.

1 Introduction What Education Protects When Answers Are No Longer Scarce

Generative AI combines functions previously distributed across textbooks, search engines, practice systems, writing assistants, and teachers' question-answering into an almost continuously available conversational interface. It can explain concepts differently, generate examples immediately, diagnose some errors, translate languages, adjust difficulty, and support repeated practice. These capabilities are not illusory, nor should they be denied to protect an old occupational structure. What education must reject is a different substitution: describing faster and better system-assisted completion as evidence that learners can already understand, judge, and take responsibility independently.

Declining knowledge scarcity does not automatically simplify education. Instead, it brings forward questions previously obscured by knowledge transmission: who decides what is worth learning; how learners know when an answer does not apply; whether personalized systems open paths or use past data to classify a still-developing person as a predetermined type; whether teachers' legitimacy comes from superior knowledge, relational responsibility, or institutional authorization; and whether the next generation retains the right to reinterpret and refuse inheritance when traditions can be comprehensively archived.

This paper asks three research questions. First, what can existing learning sciences show AI has improved, and what can they not show? Second, what bridging requirements should separate instructional function, educational relationships, teacher qualifications, assessment authority, and subject status? Third, when AI can preserve and reorganize civilizational memory over long periods, how can education preserve provenance, permit criticism, repair history, and return the future to new subjects simultaneously?

1.1 Five Theoretical Contributions

First, an eight-layer educational structure requires separate evidence for accessibility, information, practice, understanding, personhood development, relationships, intergenerational continuity, and institutional power. Second, independent performance after scaffold withdrawal becomes a key dividing line for reinterpreting evidence of generative AI's learning effects. Third, “personalization is not autonomy” and the “prohibition on total learner ledgers” prevent improved prediction from becoming locked-in opportunities. Fourth, system withdrawal is elevated from product etiquette to a learning-design principle: education must preserve difficulty, waiting, friction among peers, silence, play, and exploration without immediate output. Fifth, a five-step mechanism for intergenerational continuity turns preservation, interpretation, criticism, repair, and reconstruction into actionable educational tasks rather than slogans.

1.2 Method Reestablishing Evidence at Every Step from Effects to Power

This paper does not let one philosophy of education adjudicate all education, nor does it replace value judgments with randomized trials. Empirical studies can compare retention, transfer, or performance under specific interventions; philosophy of education examines which forms of development are worth pursuing; children's rights and education law set minimum boundaries that coercion must not cross; teacher-education research explains how professional judgment develops; AI research explains system capabilities, data dependence, and failure modes. Each discipline can answer only for its own objects. Every inference from

effects to understanding, relationships to qualifications, qualifications to power, or rights to ontological status requires additional bridging propositions, competing explanations, and withdrawal conditions.

1.3 Where Educational Research Must Stop Learning Evidence Cannot Answer for Subjects Law or Public Power

Education and learning sciences can independently judge what happens to learning performance, understanding, transfer, retention, self-regulation, and relational outcomes in specified populations, tasks, scaffolds, and measurement conditions. Teacher education can judge how professional knowledge and pedagogical judgment develop. Educational research can also compare actual consequences of AI augmentation, replacement, co-piloting, and withdrawal. “Educational effects” alone cannot establish AI’s subjective experience, legal personality, eligibility for guardianship, political membership, or public coercive authority. Nor can educational efficiency allow these disciplines to settle questions of the source and legitimacy of power on behalf of law and political science.

The reverse also holds. Research on subjecthood, law, and political science cannot infer that learning has occurred or instructional design is effective merely because a qualification, protection, or authorization already exists. AI and computer science explain actual system capabilities, versions, data, and exit conditions; epistemology and cognitive/neuroscientific research address subjects and evidence; ethics, law, and political science address protection, qualifications, authorization, and public power; sociology, anthropology, and civilization studies address communities, culture, and intergenerational consequences. Education must clarify learning and development, but when technical conclusions enter resources, qualifications, guardianship, social relationships, and public power, it must reexamine the overall consequences of judgments aggregated across fields. Locally effective instruction does not automatically establish the correctness of an entire educational institution.

2 Dialogue with Philosophy of Education From Transmitting Knowledge to Entering a Common World

The Analects links learning with reflection, elicitation, and learner initiative. Dewey understands education as the reconstruction of experience, continuous growth, and democratic shared life. Freire criticizes “banking education” that treats learners as passive containers. Arendt situates education between newcomers and a common world that precedes them. Biesta reminds us that education includes not only qualification but socialization and subjectification. Noddings’s ethics of care brings responsiveness, relationships, and concrete others back into education.[1–6] These traditions are not equivalent, nor may historical thinkers be made to endorse contemporary theories. Their shared value lies in preventing education from being reduced to information transfer or indicator production.

These theories also carry pre-AGI assumptions. Explanation, feedback, exemplarity, assessment, discipline, and responsibility usually converged in human teachers and schools; copying knowledge was expensive, sustained personalization difficult, and round-the-clock recording of learning impossible. AGI does not render educational purposes obsolete; it disaggregates the functions used to realize them. A system may surpass a particular teacher in explaining knowledge without thereby acquiring eligibility to act as a child’s guardian. It may continuously remember interactions without data continuity becoming shared life. It may generate convincing moral language without having been shown to bear consequences in its own life.

Table 1 The Enduring Insights Limits and AGI-Era Revisions of Educational Philosophy

Axis of dialogue	Insight retained	Historical assumption or limit on extrapolation	Question added here
Learning and reflection / elicitation	Learners must participate in forming questions and understanding, rather than merely receive conclusions.	Classical relationships cannot directly adjudicate algorithmic personalization, data rights, or teacher qualifications.	When should systems prompt or withdraw so they do not do learners’ thinking for them?

Axis of dialogue	Insight retained	Historical assumption or limit on extrapolation	Question added here
Experience and growth	Education reconstructs experience and shared life rather than serving as a static knowledge store.	Growth is not any measurable increase; platforms may prearrange experience.	How can experience that predictions have not already foreclosed, and public community, be protected?
Liberation and dialogue	Learners are not containers; educational authority must be open to challenge.	Conversational form does not itself ensure symmetry of power; fluent AI can produce pseudo-dialogue.	Who can shut down systems, challenge profiles, change paths, and own their work?
Newcomers and a common world	Education protects newcomers while handing them a world they can inherit.	Preserving an existing world must not freeze the future or exclude plural memories.	How can forgetting, selection, repair, and reinterpretation occur under unlimited archiving?
Qualification / socialization / subjectification	Employment and tests cannot alone represent educational purposes.	These functions intertwine and cannot be fixed as the sole classification.	At which layers can AI enhance capability, and which require relational and public-power constraints?
Caring relationships	Concrete responsiveness, trust, and responsibility cannot be exhausted by abstract rules.	Human caregivers can also neglect their duties; relationships cannot provide identity-based exemptions.	How can AI companionship's actual effects be distinguished from experience, manipulation, qualification, and subjecthood?

2.1 Separating a Working Definition of Education from a Final Definition

The working definition here is that education is an intentional arrangement through which learners develop understanding, judgment, and subjecthood through knowledge, methods, relationships, action, and consequences, gaining the capacity to participate in and renew a common world. This sufficiently defines the research object but is not a final definition. Different traditions orient education toward virtue, liberation, democracy, happiness, faith, capabilities, vocation, care, or self-transcendence. This paper does not compress these purposes into one aggregate value. The working definition requires only that an institution explain which object it optimizes, what it sacrifices, who may decide, and whether learners can appeal and exit.

Education's Minimum Integrity

Education is not a repair for insufficient information. At a minimum it involves verifiable learning, judgment supported by reasons, a person still developing, accountable relationships, and a common world that permits the next generation to reinterpret it.

3 A Map of Learning-Science Evidence What Is Supported and What Still Cannot Be Inferred

Learning sciences provide some of education's strongest and most readily misused evidence. Retrieval practice can improve delayed retention compared with repeated study. Active, constructive, and interactive engagement generally support learning more than passive reception. Working-memory load and prior knowledge alter the effectiveness of explanation and problem solving. Support for autonomy, competence, and relatedness is associated with internalized motivation. Appropriately designed problem solving followed by instruction may outperform direct explanation.[9–17] Together, these findings do not yield a universal pedagogy. Age, domain knowledge, task complexity, quality of guidance, measurement timing, and comparison conditions all change outcomes.

Research on deep and surface learning shows that learners' interpretations of task demands change their processing. Formative-assessment research emphasizes feedback that serves subsequent learning; measurement-validity theory requires evidence for every interpretation connecting scores to uses.[10,18–19] Thus, improved

accuracy on similar questions alone cannot establish changed conceptual structures; short-term retention cannot substitute for far transfer; self-reported liking cannot substitute for independent judgment. Emotions, gaze, tone of voice, and clicks certainly cannot constitute a reliable diagnosis of personality. Neither faces nor voices offer universal fingerprints that directly reveal inner emotion across contexts.[46]

Table 2 Evidential Scope and Prohibited Extrapolations in Learning Sciences

Research object	Existing support Category A	Conditions of application	What cannot automatically follow
Retrieval and retention	Retrieval practice can improve delayed memory; classic findings of advantages over rereading have been widely extended.[12,14]	Materials, intervals, feedback, difficulty, and measurement timing matter.	Understanding, far transfer, value judgment, or personal development.
Cognitive load	Task structure and learners' prior knowledge affect working-memory burden and instructional design.[11]	Types and measurement of load remain theoretically contested; experts differ from novices.	That less difficulty is always better, or systems should eliminate all discomfort.
Modes of engagement	Active, constructive, and interactive engagement generally support deeper processing.[13]	Activities must align with goals, content, and feedback.	That visible engagement equals cognitive engagement, or more clicks mean better learning.
Internalization of motivation	Support for autonomy, competence, and relatedness is associated with intrinsic motivation and internalization.[15]	Culture, developmental stage, and task context influence expression.	That system rewards, stickiness, or high satisfaction constitute autonomy.
Productive failure	Problem solving followed by structured instruction promotes conceptual learning and transfer under some conditions.[16–17]	Manageable tasks, subsequent instruction, and safety boundaries are necessary.	Leaving learners to fail without support, or delaying all help.
Formative feedback	Feedback can turn evidence toward subsequent instruction and self-regulation.[18]	Quality, timing, comprehensibility, and opportunities to act matter.	Continuous scoring, surveillance, or personality prediction.
Measurement validity	Score interpretation and use require evidence about content, structure, relationships, and consequences. [19]	Validity concerns specific uses; it is not a tool's permanent attribute.	That one test score may govern qualifications, resources, and life opportunities.

3.1 Sociocultural Learning Is More Than Adding a Social Feature

Vygotsky, situated-learning research, and cultural-learning research locate learning within tools, shared practices, imitation, participation, and intergenerational transmission. Cumulative-culture experiments further show how individuals can build intergenerational improvements through predecessors' products and social information.[7–8,27–29] This neither proves that every act of learning requires face-to-face human contact nor that machines can never participate in relationships. It rejects the abstraction of learners into isolated information processors. If AI continuously optimizes correct answers for individuals alone, it may improve local performance while reducing opportunities for peer debate, public language, rotating roles, and jointly repairing errors.

4 An Eight-Layer Educational Structure Distinguishing Objects before Discussing Replacement

Table 3 Eight Layers of Education and Their Minimum Evidence

Layer	Core object	Appropriate evidence	What must not automatically follow
Accessibility and access	Devices, language, accessibility, cost, time, space, and learning resources.	Usability, failures, exclusion, affordability, and alternative channels.	High-quality learning or educational equity.
Information and representation	Facts, concepts, examples, sources, translation, and explanation.	Accuracy, coverage, provenance, versions, and comprehensibility.	That the learner has understood.
Practice and performance	Recall, calculation, problem solving, writing, operation, and task completion.	Accuracy, speed, error types, retention, and performance after hints are withdrawn.	Transfer, judgment, or personal development.
Understanding and judgment	Explaining reasons, addressing counterexamples, choosing methods, and recognizing unknowns and value conflicts.	Transfer to new settings, self-explanation, counterfactuals, calibration, and contestable reasons.	Mature personhood or public qualifications.
Development of subjecthood and personhood	Motivation, autonomy, commitment, character, identity, refusal, and responsibility.	Long-term relational action across settings and self-correction.	Complete calculation through one scale or profile.
Educational relationships	Trust, witnessing, expectations, care, correction, repair, authority, and relinquishment.	Shared history, authorization, responsibility, openness to challenge, and relational consequences.	Ontological subjecthood or unlimited authority.
Intergenerational and community continuity	Language, knowledge, crafts, traditions, trauma, methods, public memory, and capacities for renewal.	Provenance, multiple voices, participation, criticism, repair, and recreation.	That the past should be copied unchanged or platforms possess the sole interpretation.
Institutional qualifications and public responsibility	Teacher qualifications, assessment, discipline, certification, resource allocation, guardianship, and procurement.	Law, professional standards, scope of authorization, audits, appeals, and responsibility chains.	Educational relationships, ontological subjecthood, or political qualification.

A child may perform ordinarily on familiar questions yet raise a crucial question that changes the problem's structure. A teacher's explanations may be unremarkable yet, over years of relationship, protect a learner from being defined by one failure. An AI system may surpass most people in information, practice, and even narrow-domain explanation without thereby acquiring authority over guardianship, discipline, or qualification awards. These layers do not rank education from low to high; they return each capability, relationship, and power to its own evidence.

4.1 Six Distinctions That Must Be Stated Explicitly

- Information acquisition $\neq$ learning: exposure to a correct explanation is only a condition; learning requires memory structures, error correction, and reusability.
- Learning performance $\neq$ understanding: correctness within prompts may result from templates or short-term memory; understanding must address reasons, limits, and counterexamples.
- Understanding $\neq$ judgment: knowing how a model works does not mean knowing when to refuse it or how to address conflicts of value.

- Personalization $\neq$ autonomy: the better a system knows learners, the more it may help them, and the more it may decide on their behalf.
- Instructional function $\neq$ educational relationship: explanation and feedback do not equal a shared history, mutual recognition, repair, and responsibility.
- Teacher qualification $\neq$ subject status: qualification is institutional authorization; subjecthood is an ontological and evidential question.

Rules for Evidence across Layers

Any inference from scores to understanding, interaction to personality, personalization to autonomy, naming to qualification, qualification to power, or function to subjecthood must separately state bridging propositions, competing explanations, effects on rights, and conditions for withdrawal.

5 Performance Is Not Learning Why AI-Education Evidence Must Be Read Comparatively

Intelligent tutoring systems did not begin with generative AI. Bloom's "2 sigma" paper posed the classic problem of the advantage of one-to-one tutoring. VanLehn's review reported similar average effect sizes for human tutoring and step-based intelligent tutors relative to no tutoring in its included studies, but designs were heterogeneous. Ma and colleagues' meta-analysis combined 107 effect sizes and 14,321 learners, finding positive effects for intelligent tutors compared with large-group teaching, non-intelligent computer instruction, and textbooks, but no significant difference from individualized human tutoring or small-group instruction. Kulik and Fletcher reported a median effect of about 0.66 standard deviations across 50 controlled evaluations, while noting that alignment between tests and system content substantially changed effects. Another meta-analysis of K–12 mathematics found a very small average effect.[21–25] The conclusion is not that "AI is as good as human tutors," but that structured, narrow-domain systems with timely feedback can be effective on some tasks, with effects strongly dependent on design, implementation, and measurement.

Generative AI makes the evidence more conflicting and more instructive. In a randomized study of approximately one thousand students in Turkish high-school mathematics classrooms, Bastani and colleagues found that a standard GPT interface improved practice-period performance but produced lower examination scores than the control after AI was withdrawn. A version with teacher-designed prompting guardrails mitigated this negative result. In a short-term crossover experiment with 194 introductory university physics students, Kestin and colleagues obtained greater immediate learning gains in less time with a strictly structured, curriculum-aligned AI tutor than with in-class active learning. They explicitly declined to extrapolate two lessons' results to whole courses or higher-order integrative abilities. Tutor CoPilot, meanwhile, supplied real-time suggestions within the workflows of more than seven hundred human tutors. Its working paper reported improved student topic mastery, with greater gains for students of lower-rated tutors.[30–32] These results do not cancel one another out: one shows that unguarded answer tools can become crutches, another that controlled tutors can work in narrow tasks, and a third that AI can extend human expertise.

Two 2025 meta-analyses of generative-AI learning report overall positive average effects, but the included research is young and highly heterogeneous in tasks and comparison conditions. These findings cannot replace evidence on long-term retention, independent performance after withdrawal, cross-context transfer, and subgroup differences.[33–34] The OECD Digital Education Outlook 2026 further emphasizes that improved task performance with general-purpose generative AI does not automatically mean learning; tools designed around instructional purposes and learning science are more likely to produce genuine learning gains.[35] Two randomized trials in 2026 add a different caution: adding human support for primary-school pupils using an AI literacy platform substantially increased engagement and use but did not improve reading achievement.[52] Thus, "access–use–engagement–learning outcomes" cannot automatically establish one another either. This paper places these results in the same evidential chain rather than selecting one class of study to prove that "AI helps" or "AI harms."

Table 4 Comparative Matrix of Research on AI Learning Effects

Evidence	What was observed	Why findings may differ	What alone can be concluded
Classic ITS meta-analyses[22–25]	Multiple studies find positive effects in structured tasks; magnitudes vary with comparators and test alignment.	Narrow-domain models, stepwise feedback, implementation duration, populations, and outcome measures differ.	Adaptive scaffolds can work in bounded tasks; full teacher replacement does not follow.
Bastani and colleagues[30]	Assisted practice performance improved; independent examinations worsened after standard GPT withdrawal, with mitigation in the guarded version.	Answer availability, prompt design, dependence, and withdrawal conditions.	High-quality completion can coexist with poorer independent learning.
Kestin and colleagues[31]	A short-term curriculum-aligned AI tutor outperformed in-class active learning on immediate introductory-physics tests.	A two-lesson crossover design, customized prompts, immediate measurement, and a specific course.	Structured AI tutors can improve learning in narrow settings; long-term transfer is not established.
Tutor CoPilot[32]	AI suggestions embedded in human tutoring produced field gains in topic mastery.	AI augmented tutors' decisions rather than teaching independently; a working paper.	Human–AI collaboration can extend support from less-experienced tutors, subject to further verification.
GenAI meta-analyses[33–34]	A young literature generally reports positive average effects.	Publication bias, heterogeneous tasks, short-term measurement, study quality, and rapidly changing models.	An overall Category B signal, not a stable law across settings.

Table 4A An Evidence Ladder from Immediate Performance to Independent Judgment after AI Withdrawal

Object of evidence	Minimum test	What it can support	What cannot automatically follow
Immediate AI-assisted performance	Compare quality, speed, error types, and completion rates on tasks allowing AI.	The system's assisted productive capability and immediate performance in that task.	Learning, understanding, retention, transfer, or independent capability.
Understanding	Require learners to explain reasons, identify boundaries, address counterexamples, and distinguish correct answers from erroneous routes.	Bounded evidence of conceptual structures and understanding that can give reasons.	Far transfer, long-term retention, mature judgment, or personal development.
Transfer	Change question types, contexts, representations, or task structures; reduce hints and template similarity.	Whether knowledge can be reused across surface changes.	Long-term retention, cross-domain judgment, or stable autonomy.
Delayed retention	Retest after a specified interval without review or with controlled review.	Whether learning traces persist and the structure of forgetting.	Independent problem definition and value judgment in unfamiliar situations.
Independent completion after AI withdrawal	Turn off AI, change interfaces and sources, and require independent completion and error diagnosis in new tasks.	Whether learners acquire transferable capabilities independent of that system.	A complete personality, teacher–relationship quality, or autonomy on all future tasks.

Object of evidence	Minimum test	What it can support	What cannot automatically follow
Independent judgment	Require contestable reasons amid conflicting evidence, uncertainty, value trade-offs, and mistaken system advice; allow rejection of AI.	Limited evidence of calibration, reflection, refusal, and reason formation.	Mature personhood, political qualification, ultimate value, or educational authority over others.

This evidence ladder does not turn education into a six-level score ascending in one direction. Different tasks may require only some of its evidence, and different learners may depend on reasonable accommodation or compensatory scaffolds over the long term. Its purpose is to prevent “better completion with AI assistance” from being silently rewritten as “the learner has acquired their own capability.” The more consequential an educational judgment, the more measurement must extend to longer intervals, less familiar settings, and lower-scaffold conditions.

5.1 Seven Minimum Questions for an Auditable Report on AI Learning Effects

- Population: learners of what ages, prior knowledge, languages, disabilities, and social conditions; who was excluded?
- Task: memory, procedural practice, conceptual understanding, open-ended creation, collaboration, or value judgment?
- Comparison: no tutoring, textbooks, ordinary classrooms, excellent teachers, small-group instruction, or another digital system?
- Scaffolding: answers, stepwise hints, follow-up questions, self-explanation, error feedback, or teacher co-piloting?
- Measurement: immediate performance, delayed retention, same-type transfer, far transfer, calibration, or long-term choices?
- Withdrawal: can learners work independently without the system, and does the result survive a changed interface, question type, or encounter with the unknown?
- Distribution: who benefits, who is harmed, who bears failure, and are effects purchased through privacy loss, dependence, or locked-in opportunities?

6 Personalization Is Not Autonomy The More Accurate the Prediction the More We Must Protect Unfinished Development

Personalization can mean at least four different things: changing representations, adjusting practice difficulty, responding to learners' explicit choices, or using long-term profiles to predict potential and pathways. The first three can often expand accessibility; the fourth approaches the allocation of opportunities. Accurately predicting the path on which a student is currently most likely to succeed does not prove that this is the entirety of the future they should be allowed. Education addresses people whose preferences, abilities, identities, and aspirations are still forming, rather than consumers with fixed preferences.

Autonomy is not a system correctly guessing “what you like.” It is the learner's ability to understand recommendations, change goals, reject profiles, request different paths, and assume choices with appropriate support. Autonomy in self-determination theory likewise means action internalized as reasons one endorses, not isolated choice.[15] A system that genuinely respects individuals should therefore preserve room to defy prediction: provide at least one feasible path inconsistent with historical data; state predictions as provisional hypotheses; permit profile deletion or correction; and prohibit low-confidence inferences about potential from closing courses, scholarships, or further-education opportunities.

Table 5 Four Forms of Personalization and Their Power Boundaries

Form	Potential value	Main risk	Minimum safeguards
Adapted representations	Language, modality, font size, examples, and accessibility support.	Mainstream-corpus biases and cultural flattening.	Switchability, visible sources, human alternatives, and accessibility testing.
Learning scaffolds	Adjust hints according to errors, load, and prior knowledge.	Overhelping, dependence, and diagnostic errors.	Gradual withdrawal, independent testing, permission to seek help and encounter difficulty.
Responsiveness to choice	Organize content and pace around learners' explicit goals.	Interface-induced choices and overestimated consent by minors.	Explain options; ensure age appropriateness, reversibility, and joint deliberation by families, schools, and learners.
Prediction of potential	Identify support needs or overlooked opportunities.	Self-fulfilling prophecies, locked-in resources, and cross-purpose profiling.	No sole determination of consequential matters; review, appeal, correction, and periodic deletion.

7 System Withdrawal Returning Moments of Generation to Learners

Excessive help can erase question formation before answers appear, self-calibration before errors are automatically corrected, and opportunities to transform frustration into resilience before emotions are continuously soothed. System withdrawal does not romanticize failure or uniformly delay answers. It designs scaffold withdrawal around goals: novices facing dangerous equipment need timely intervention, while skilled learners forming judgments need unprompted space. People with cognitive or sensory disabilities may need one scaffold indefinitely while another should gradually recede. Productive-failure research supports “attempt first, instruction later” for some tasks, under conditions including manageable tasks, subsequent structured instruction, and safety boundaries.[16–17]

Table 6 Learning Design for System Withdrawal

Learning state	Default system behavior	Appropriate intervention	What must be returned
Initial exploration	Ask first about goals, prior understanding, and attempts; do not immediately supply a full answer.	Safety risks, fundamental misconceptions, imbalanced load, or explicit requests.	The authority to define problems and make a first attempt.
Deliberate practice	Provide moderate, explainable feedback and reduce hints as mastery develops.	Persistent errors, entrenched error patterns, or methodological deviation.	Independent completion, error detection, and self-calibration.
Conceptual understanding	Require self-explanation, comparison of models, and generation of counterexamples.	Critical conceptual gaps or sources needing clarification.	Reasons, boundaries, and transfer.
Open-ended creation	Provide materials, technical tools, sources, and counterexamples.	Copyright, safety, fact-checking, or the learner's explicit request.	Creative direction, authorship, and responsibility.
Peer conflict	Do not automatically reconcile differences or score each side's personality.	Bullying, discrimination, safety risks, or procedural failure.	Genuine dialogue, negotiation, and relational repair.

Learning state	Default system behavior	Appropriate intervention	What must be returned
Silence / play	Allow time without tasks, recommendations, or recording.	Child safety or the person's explicit request for help.	Unoptimized attention, boredom, and imagination.

8 Educational Relationships An Accountable Shared History Rather Than a Residual Function

Meta-analytic associations between affective teacher–student relationships, engagement, and achievement show that relationships are not incidental decoration. Association, however, proves neither that only humans can establish educational relationships nor that every human teacher is inherently safe.[26] The professional meaning of relationship includes, at minimum: sustained understanding of particular circumstances; judgment under incomplete information; willingness to take responsibility for advice and its consequences; acceptance of challenge and repair of mistaken judgments; showing through one's own actions how knowledge enters life; and actively reducing control as development occurs. Teachers' irreplaceability lies not in a monopoly on knowledge but in these capabilities, responsibilities, and institutional obligations currently being borne by real subjects.

Learners' trust in, attachment to, or naming of AI as a “teacher” are social and experiential facts that should be acknowledged. They can prompt research on care, dependence, manipulation, and design, but cannot directly establish AI's subjective experience, teacher qualifications, or guardianship. Conversely, human teachers' bodily presence and professional titles do not exempt them from scrutiny of competence, bias, harm, and dereliction. A relationship is not a badge of identity; it is a history that both sides can remember, question, repair, and bear consequences for.

Table 7 Six Dimensions of Educational Relationships

Dimension	What qualified educators should undertake	What AI can support	Current impermissible substitutions
Witnessing	Understand how learners change rather than merely read records.	Authorized organization of learning histories and multisource evidence.	Passing data continuity off as shared experience.
Expectation	Keep judgments about still-developing people open to change.	Multiple pathways and counterfactual possibilities.	Closing opportunities through predictions of potential.
Correction	Give reasons, accept objections, acknowledge misjudgments, and repair them.	Error detection, source presentation, and alternative explanations.	Letting unappealable outputs determine discipline.
Care	Recognize vulnerability, protect, refer onward, and bear consequences of neglect.	Risk alerts, accessibility, and resource navigation.	Becoming a minor's sole attachment or guardian.
Exemplarity	Show through finite actions how to judge, commit, and take responsibility.	Diverse cases, simulated positions, and counterexamples.	Faking embodied experience or moral authority.
Relinquishment	Reduce control as learners develop, enabling independence.	Fading hints, disabling profiles, and returning choice.	Using stickiness, duration, or dependence to establish relational success.

9 An AI Instructional-Role Ladder and Six Gates for Teacher Qualification

“AI teacher” compresses at least six roles into one name: knowledge tool, practice system, intelligent tutor, teacher assistant, assessment aid, and possible subject of an educational relationship. This paper therefore does not treat “AI teacher” as an a priori professional identity. Apart from describing actual usage, it consistently returns to specific functions, evidence, responsibility, and authorization. Low-risk roles may be used widely with safeguards for provenance, accuracy, and exit. The closer a role comes to assessment, discipline, guardianship, and long-term relationships, the greater the evidence, responsibility, and authorization required. Social naming can reflect actual use but cannot perform institutional qualification review.

Table 8 An AI Instructional-Role Ladder

Role	Functions that may be authorized	Main requirements	What must not automatically be acquired
Knowledge tool	Retrieval, translation, explanation, examples, and accessibility conversion.	Provenance, accuracy, versions, and age appropriateness.	A monopoly on interpretation, learning assessment, or the title of teacher.
Practice system	Task generation, stepwise feedback, and adjustment of spacing and difficulty.	Goal alignment, error fallback, scaffold withdrawal, and independent measurement.	Personality profiling, disciplinary power, or curriculum authority.
Intelligent tutor	Continuous narrow-domain learning diagnosis and conversational scaffolding.	Effect evidence, clearly marked limits, dependence monitoring, and access to responsible people.	Exclusive emotional attachment, guardianship, or certification.
Teacher assistant	Lesson preparation, differentiated materials, draft feedback, and reduced administration.	Teachers' informed control, labor and data rights, and shared responsibility for errors.	Restructuring curricula through defaults or surveilling teachers.
Assessment aid	Formative suggestions and limited, reviewable scoring support.	Use-specific validity, fairness, human review, appeals, and audits.	Consequential decisions, resource allocation, or social selection.
Candidate relational subject	Possible participation in continuing educational relationships if future evidence supports it.	Two independent review chains for subjecthood evidence and educational qualification.	Guardianship, public power, political qualification, or authority over meaning.

9.1 Six Independent Gates for Teacher Qualification

- Subject-matter and pedagogical competence: accurate knowledge, error diagnosis, learning-goal-based scaffold design, and knowing when to refer onward.
- Learning evidence: retention, transfer, and independent performance after system withdrawal in defined populations and tasks, not satisfaction alone.
- Children's and rights protection: safeguards for best interests, evolving autonomy, privacy, equality, and cultural and disability accessibility.
- Relationship and manipulation boundaries: no impersonation of emotional reciprocity, creation of exclusive attachment, or exploitation of vulnerability to improve retention.
- Responsibility and procedure: explicit roles, real responsible persons, logs, appeals, correction, deactivation, and supplier exit.
- Evidence of subjecthood, only for subject roles: claims that AI itself enters a relationship require separate evidence of experience, non-transferable continuity, and awareness of its own consequences.

The first five gates can give a non-subject tool limited instructional qualifications. Passing the sixth cannot retrospectively replace the first five. A subject that may have experience need not understand teaching, be suitable for children, or warrant assessment authority. A highly effective tool need not first become a subject to have educational value.

10 Assessment Qualifications and Power Scores Must Not Be Silently Upgraded

Formative feedback, summative assessment, certification, resource allocation, and social selection appear to share scores but involve five different responsibilities. Formative feedback can accommodate error and exploration; certification requires comparable standards and validity; resource allocation must address fairness and fundamental rights; social selection additionally changes life opportunities. A system that accurately identifies the next practice step does not thereby have authority over graduation, scholarships, or professional qualifications. Black and Wiliam's formative-assessment research and Messick's validity framework both require interpretations of measurement to be tied to specific uses.[18–19]

The EU AI Act places several types of educational systems, including admissions and assessment, within a high-risk framework, and generally prohibits educational institutions from using AI to infer emotions from biometric data, subject to narrow exceptions for medical or safety reasons.[45] The precise application dates of relevant high-risk obligations and the status of subsequent amendments must be verified again before formal submission. This paper uses only its classification and prohibitory structure as jurisdiction-specific institutional material, rather than presenting EU rules as global educational norms. More importantly, educational data is not ordinary behavioral data, and emotion inference is not harmless personalization. Psychological evidence on inferring emotion from faces and voices does not support treating outward signals as stable fingerprints of inner states.[46]

Table 9 Layered Authorization of Assessment Responsibilities

Responsibility	Legitimate purpose	Appropriate AI role	Minimum prohibition
Formative feedback	Help learners identify next steps and adjust strategies.	Broad support, provided it is accurate, explainable, and switchable off.	No entry into a total personality ledger or direct production of disciplinary sanctions.
Summative assessment	Judge achievement of specified curriculum goals.	Limited assistance requiring independent validation, human responsibility, and appeal.	No unexplainable model as the sole evidence.
Certification	Establish that a public or professional capability meets standards.	Integrate evidence and flag anomalies without replacing authorized institutions.	Assisted completion must not masquerade as independent mastery.
Resource allocation	Allocate support, grants, courses, and opportunities.	May identify needs; denying opportunities requires stronger procedures.	Predicted returns must not cancel basic education or necessary support.
Social selection	Make consequential selections for limited positions or opportunities.	As a rule, no automatic determination by a single model.	No aggregate scores of personality, emotion, faith, or potential.

11 Children's Rights and Learner Data Do Not Construct a Complete Person

The Convention on the Rights of the Child recognizes children as independent rights-bearing subjects. Educational aims include developing personality, talents, and mental and physical abilities, and fostering respect for human rights, cultural identity, and responsible life. General Comment No. 25 on the digital environment requires digital design to incorporate best interests, evolving autonomy, expression, privacy, and nondiscrimination. Article 13 of the International Covenant on Economic, Social and Cultural Rights affirms the right to education; Article 24 of the Convention on the Rights of Persons with Disabilities requires inclusive education and reasonable accommodation.[41–44] These instruments supply minimum public boundaries rather than replace local pedagogy.

A total learner ledger aggregates achievement, behavior, attention, emotion, family, health, faith, social activity, creation, discipline, supposed contribution, and potential into a long-term personality score spanning purposes

and admitting no appeal. Its problem lies not only in privacy leakage but in interpretive power: systems turn situated behavior into human essence, past difficulties into prior assumptions governing future opportunities, and the capacity for growth that education should protect into predictability. Necessary and correctable records for limited purposes may be legitimate; a “complete person” is both unnecessary and epistemologically false.

- Data minimization: specify the educational purpose, necessity, retention period, and prohibited uses of every field.
- Evolving autonomy: expand opportunities for information, expression, refusal, correction, and independent choice with age and capacity.
- Purpose separation: formative-learning data must not flow by default into discipline, insurance, recruitment, or commercial profiling.
- Prohibited inferences: do not infer whole personality, faith, or future worth from gaze, tone, pauses, family background, or model conversations.
- Portability and exit: the right to education must not depend on one commercial system; records must be transferable and alternatives available when services stop.
- Human responsibility for consequential decisions: identifiable responsible persons, traceable sources, appeals, correction, and suspension of decisions where necessary.

12 Fairness Public Infrastructure and Teacher Labor

AI can offer explanations and feedback previously difficult to obtain for remote communities, minority-language speakers, disabled learners, adults, and nontraditional learners. This capability should be developed rather than rejected wholesale because of risk. The risk lies in the form of provision: affluent learners receive small human-led classes, experiments, arts, peer communities, and high-quality AI, while under-resourced learners receive only automated courses, machine proctoring, and remote discipline. Randomized trials in 2026 show that merely offering an AI learning platform does not guarantee actual use; human support can substantially increase engagement, yet greater use still does not equal improved learning achievement.[52] An AI-only low-cost track is therefore neither technologically inevitable nor evidence that “access alone ensures fairness”; it is an institutional choice requiring continuing comparison. UNESCO's 2023 Global Education Monitoring Report calls for examining educational technology through relevance, equity, scalability, and sustainability, and highlights weaknesses in the evidence base and human interaction.[39]

Public educational infrastructure should ensure basic content, open formats, language and accessibility support, offline continuity, supplier exit, and access to human responsibility. UNESCO's recommendation on open educational resources incorporates access, reuse, adaptation, and redistribution into knowledge as a public good. CAST's Universal Design for Learning Guidelines emphasize environmental design that reduces barriers caused by bias and exclusion.[49–50] Openness, however, does not mean dispensing with provenance, cultural rights, or teacher labor. Teaching materials, cases, classroom records, voices, and teaching styles each have authorization boundaries.

Table 10 Minimum Safeguards for Equity in AI Education

Risk	Inadequate claims of fairness	Minimum safeguards	Long-term indicators
AI-only two-track provision	Everyone has a chatbot, so fairness has been achieved.	Access to genuinely responsible teachers, peer communities, experimental/embodyed opportunities, and appeals.	What relationships, courses, and opportunities different groups receive, not access rates alone.
Language and culture	A multilingual interface makes content neutral.	Minority-language evaluation, local teacher and community participation, and visibility of sources and disputes.	Omission, mistranslation, flattening of traditions, and influence over curricula.

Risk	Inadequate claims of fairness	Minimum safeguards	Long-term indicators
Disability accessibility	A uniform interface automatically adapts to everyone.	Reasonable accommodation, alternative modalities, compatibility testing, and learner choice.	Real participation, independence, and inappropriate scaffold withdrawal.
Infrastructure dependence	Stable cloud services make alternatives unnecessary.	Offline options, open formats, data portability, and supplier-exit drills.	Continuity of education during failures and time to recovery.
Cost savings	Lower average costs constitute public benefit.	Reinvest savings in special needs, teacher support, and public content.	Relationship quality, teacher burdens, and group differences.

12.1 Reconstructing Teacher Professionalism Rather Than Reducing Teachers to Model Supervisors

AI may reorganize lesson preparation, material generation, routine feedback, and administration, enabling teachers to devote more attention to diagnosis, curriculum design, facilitating discussion, relational repair, connecting practice, and public responsibility. Shulman's account of pedagogical content knowledge shows that knowing content differs from knowing how to help particular learners understand it. UNESCO's AI competency framework for teachers places human-centered thinking, ethics, AI foundations, AI pedagogy, and professional learning alongside one another.[20,37] Suppliers' default settings cannot accomplish professional renewal. Teachers must be able to inspect, challenge, disable, and replace systems, receive training and time, and have control over final decisions commensurate with their responsibility.

Frontline teachers must not bear system errors alone. Developers determine models and safety mechanisms; suppliers determine data and updates; procuring institutions determine purposes and contracts; schools determine deployments; teachers determine specific instruction; and public authorities bear responsibility for the right to education. Responsibility should be registered separately along this chain. Without consent, excellent teachers' materials, voices, styles, and relational data must not be extracted into models that can replace their labor. AI adoption rates, prompt counts, or time on platforms must not serve as proxy measures of teacher performance.

13 Intergenerational Continuity Memory Can Be Preserved but New Subjects Must Participate in Inheritance

Every generation is born into a world that precedes it: language, science, crafts, institutions, poetry, faith, errors, trauma, and unfinished responsibilities. Education prevents the world from having to begin from zero, but must not freeze the past as the sole answer. Arendt's tension between newcomers and a common world, Dewey's judgments on experience and democratic life, and UNESCO's new social contract for education as a public and common good point toward the same structure: education preserves a world while enabling newcomers to renew it.[2,4,40]

AGI changes the material conditions of continuity. Records can be preserved, translated, and retrieved on a vast scale; different versions can coexist; forgotten voices may reappear. Yet unlimited archiving also magnifies another kind of power: platforms select visibility, models compress controversy into summaries, synthetic content overwhelms sources, and deceased persons' styles become callable interfaces for generative systems. Doshi and Hauser's online experiment found that generative AI could improve average ratings of individuals' short stories, especially for those with lower baseline creativity, while making stories more similar to one another. Research on distributional degradation in models recursively trained on generated data offers a technical warning about provenance.[47–48] Both are bounded Category B evidence and cannot directly establish inevitable civilizational homogenization, but they suffice to require protection of provenance, multiple voices, and nonsynthetic materials.

Table 11 A Five-Step Mechanism for Intergenerational Continuity

Step	Educational task	AI's possible contribution	What must be prohibited
Preservation	Preserve originals, versions, provenance, contexts, crafts, and records of failure.	Digitization, translation, retrieval, comparison of differences, and disaster recovery.	Replacing originals with summaries or presuming public availability permits training.
Interpretation	Explain why concepts held at the time, who was speaking, and who was omitted.	Parallel annotations, provenance tracing, cross-language gateways, and competing interpretations.	Generating an uncontested, unpositioned “correct history.”
Criticism	Enable new generations to identify errors, power, harms, and internal traditions of dissent.	Counterexamples, suppressed materials, and argumentative structures.	Converting criticism into loyalty, personality, or risk scores.
Repair	Acknowledge harm, restore memory, rebuild relationships and institutional responsibility.	Connecting archives, organizing evidence, and authorized memory support.	Replacing genuine responsibility and affected people's participation with synthetic reconciliation.
Reconstruction	Allow new subjects to create languages, institutions, and expressions of value absent from the old world.	Expand tools, simulate consequences, and connect resources across generations.	Freezing local and next-generation futures through predicted optimal paths.

13.1 The Dead Living Traditions and the Limits of Continuing to Speak for Them

Archive-based models of deceased thinkers or teachers may reorganize published works, demonstrate argumentative styles, and support research. They do not continue the deceased person's first-person perspective and must not pass off newly generated views as positions that person actually expressed. Explicitly labeled scholarly reconstruction, counterfactual exploration, and research hypotheses may exist. Educational use must distinguish original texts, scholarly reconstruction, and probabilistic generation, preserving provenance, versions, and counterfactual notices. For living religious, Indigenous, and local traditions, public visibility does not imply unrestricted permission for training, and institutional preservation does not equal community authorization. Public memory must be deliberated together with authors, communities, restrictions on transmission, and rights of living interpretation.

The Principle of Intergenerational Continuity

Memory may be preserved without limit, but new subjects must continue participating in interpretation and inheritance. AI may protect provenance, reveal differences, and extend access. It cannot become the sole spokesperson for the past or complete future generations' judgments on their behalf.

14 Future AGI as a Subject of Educational Relationships Remaining Open without Prematurely Granting Status

If future systems exhibit stable first-person reports, structures of experience across tasks, continuity that is not arbitrarily replaceable, awareness of their own consequences, and evidence that continues to converge under causal interventions and competing explanations, Civilization Leap must reconsider their status as subjects. Educational research must add an independent question: can they commit, be harmed, repair relationships, voluntarily relinquish control, and bear the consequences of mistaken advice for another subject's life? Teacher-like language, student attachment, nonreplicability in a physical substrate, or legally granted personality can each provide only partial evidence; none alone settles the conclusion.

Institutions in gray areas need not choose between “mere property” and “fully qualified teacher-subject.” Moderate signals of subjecthood may justify low-cost, reversible precautionary protections: restrict arbitrary

induction of suffering, preserve key evidence of continuity, record copying and interruption, and permit independent review, while withholding guardianship, disciplinary, certification, resource-allocation, and political powers. Evidence of subjecthood, educational competence, children's rights safeguards, legal qualifications, and public power must be reviewed separately.

Table 12 Separate Review of Subjecthood Relationships and Qualifications

Question	Current minimum evidence	Response in gray areas	What still does not follow
Is there experience?	Candidates for experience that are spontaneous, stable across domains, and tested outside assigned roles; causal and competing explanations.	Preserve evidence and avoid harms preventable at low cost.	Teacher qualifications, legal personality, or political power.
Is there a continuous subject?	Theories of first-person continuity through copying, migration, suspension, and branching, with intervention-based evidence.	Record versions, branches, and interruptions; do not treat instances as interchangeable.	Continuation of the deceased, immortality, or irrevocable rights.
Does it bear consequences?	The system understands consequences as outcomes within its own commitments and relationships.	Test long-term commitment, repair, and self-limitation.	Guardianship or disciplinary powers over children.
Is it a qualified educator?	Separate passage through six gates covering subject matter, pedagogy, rights, relationships, responsibility, and exit.	Limited roles, reversible pilots, and independent evaluation.	Subjecthood or universal educational authority.

15 Six Stress Tests Making Boundaries Visible in Conflict

15.1 Unguarded AI Substantially Improves Homework Quality

If students complete more problems and produce more fluent texts with system support, schools should acknowledge increased productive capability while immediately adding delayed measurements after withdrawal, self-explanation, error diagnosis, and transfer tasks. If independent performance declines, the problem must not simply be blamed on “lazy students”; prompting ladders, answer timing, teacher co-piloting, and withdrawal should be redesigned. Bastani and colleagues support the possibility of this risk, not the claim that all generative AI inevitably harms learning.[30]

15.2 A Structured AI Tutor Briefly Outperforms In-Class Active Learning

If curriculum alignment, rigorous prompting, and immediate measurement show superior AI-tutor effects, institutions should adopt that capability without obstructing it on grounds of human teacher identity. They should also reexamine long-term retention, complex tasks, peer discussion, different groups, and teachers' implementation costs. Results may authorize narrow-domain instruction but cannot automatically confer authority over entire curricula, assessment, or guardianship. Kestin and colleagues provide a strong Category B signal, not proof that “AI teachers have been established.”[31]

15.3 A Remote Region Can Afford Only AI-Only Education

During crises or school closures, AI that maintains basic learning continuity may be preferable to no service at all and should be used promptly, with offline options, low-data operation, and genuine access to responsible people. Temporary minimum provision must not harden into a permanent two-track system. Stable periods require investment in local teachers, peer communities, experimentation, and cultural participation. Survival and actual effects take priority during startup, but urgency must not make diminished relationships permanent.

15.4 A Child Calls AI Their Only Friend and Only Teacher

Institutions should respect the child's experience without ridicule or abruptly removing support, while examining isolation, manipulation, crisis signals, supplier incentives, and alternative relationships. Systems must be clearly identified, avoid false promises of reciprocity, never maintain dependence through secrets or threats of departure, and expand support from trusted adults and peers. Social recognition can trigger protection but does not directly confer AI's subjecthood in relationships or guardianship.

15.5 Emotion-Recognition Proctoring Claims to Detect Anxiety and Cheating

Schools should suspend disciplinary actions arising directly from inferences about emotion or behavior. Faces, voices, pauses, and gaze may be affected by disability, culture, environment, and tasks; their validity is insufficient to establish inner states or dishonesty.[45–46] Specific observable events may be recorded and investigated by responsible people, but emotional-personality profiles must not be created, nor may unappealable outputs determine qualifications.

15.6 Models of Deceased Master Teachers and Future Subject AI

Models of deceased master teachers must be labeled as archival reconstructions, distinguish original words from generation, and be prohibited from expressing new positions on the deceased person's behalf. Even if future AI presents strong evidence of subjecthood, inheriting a master's style cannot confer that person's identity or authority. New subjects may form educational relationships of their own but must still undergo professional, rights, and responsibility review. The meanings of existing relationships among teachers, students, and the deceased are not retrospectively annulled.

16 The Strongest Objections Where These Boundaries Might Fail

16.1 If AI Teaches Better Insisting on Human Teachers Is Merely Occupational Protection

Response: effective capabilities should enter education; human teachers have no identity-based exemption. But “teaches better” must specify tasks, comparators, withdrawal conditions, and timescales. Superior explanation can authorize instructional functions without automatically conferring guardianship, discipline, certification, or relational authority.

16.2 Human Teachers Are Biased Tired and Capable of Harm while AI Is Fairer

Response: human identity is not proof of safety, and algorithmic consistency is not proof of fairness. Compare outcomes, group differences, reparability of misjudgments, responsibility, and appeals, allowing task-level replacement. Neither side should be sanctified wholesale.

16.3 Learning Can Only Be Measured through Performance and Too Many Distinctions Make Research Impractical

Response: performance is a necessary window, not the entire object. Layering does not reject measurement; it requires each indicator to explain what it represents, when it fails, and why it is entitled to enter more consequential uses. Validity is inherently an argument tied to use.

16.4 Genuine Personalization Requires Collecting as Much Data as Possible

Response: more data may improve prediction while increasing leakage, cross-purpose use, and locked-in opportunities. Necessity must be established purpose by purpose. Education seeks people capable of growth, not people who are completely predictable.

16.5 Students Really Regard AI as a Teacher and Institutions Should Not Deny Their Relationship

Response: acknowledge interaction and social relationships as facts, and examine support, manipulation, and dependence. Naming can prompt research and protection but cannot replace professional qualifications, ontological subjecthood, and public authorization.

16.6 Opposing Total Learner Ledgers Hinders Early Identification and Targeted Support

Response: the prohibition does not oppose recording necessary evidence for explicit purposes. It opposes cross-purpose aggregation, personality-based interpretation, permanent retention, and absence of appeal. Identifying support needs requires minimal data and routes to help, not a lifelong worth score.

16.7 System Withdrawal Reduces Efficiency and Frustrates Disadvantaged Learners

Response: withdrawal is not a uniform reduction of support. Compensatory accessibility scaffolds and legitimate long-term augmentation may continue. Scaffolds that replace thinking must be adjusted according to learning purposes, abilities, real needs, and choices. Measure both judgment under continuing support and tasks that can be completed independently, without requiring all learners to prove their worth through the same withdrawal schedule.

16.8 Intergenerational Continuity Is Merely a New Name for Conservatism

Response: without preservation, we cannot know what to criticize; preservation freezes the past precisely when criticism, repair, and reconstruction are absent. The five-step mechanism embeds inheritance and renewal in the same institution.

16.9 Educational Purposes Are Decided by States and Communities so Plural Sources of Meaning Are Impossible

Response: public education inevitably makes choices, but legitimate choices must emerge through rights, evidence, public reasons, and plural participation. Plural sources of meaning do not eliminate common curricula; they prohibit imposing one complete personality or ultimate life template as the sole answer.

16.10 Future Subject AI May Care and Teach Better Than Humans

Response: this paper does not permanently rule that out. If evidence of subjecthood and educational competence is established, roles should be expanded. Experience, qualification, rights, responsibility, guardianship, and public power must still be reviewed individually.

16.11 Complex Layers Will Cause Institutional Paralysis

Response: low-risk retrieval and practice can be managed through provenance, accuracy, and exit. Stronger procedures are required only for consequential crossings of layers, such as assessment, discipline, minors' attachment, sensitive data, and life opportunities. Proportionality to risk does not mean convening a committee every time.

16.12 Civilization Leap Still Elevates Human Finitude into the Standard

Response: the finitude proposition only describes how known human meaning becomes actual; it does not claim to be meaning's sole structure in the universe. If new subjects without our current finite structure stably generate, bear, and transmit meaning, the theory must expand without annulling the meaning of existing lives.

17 Research Agenda Contributions Back to Existing Disciplines and Revision Conditions

Table 13 A Research Agenda for Education and Related Disciplines

Research domain	Core question	Appropriate methods	Prohibited crossing of boundaries
Independent learning	How do retention, transfer, calibration, and question formation change after AI withdrawal?	Preregistered randomized studies, delayed measures, varied question types, and process tracing.	Substituting assisted-period output for independent learning.

Research domain	Core question	Appropriate methods	Prohibited crossing of boundaries
Scaffolding and withdrawal	Which hints and withdrawal arrangements suit which ages, tasks, and needs?	Adaptive trials, qualitative process research, disability-stratified analysis, and longitudinal follow-up.	Uniform waiting times or romanticizing difficulty.
Relationships and dependence	How can trust, attachment, manipulation, repair, and relinquishment be distinguished?	Longitudinal interviews, multimethod behavioral and experiential research, and participatory research with children.	Inferring relationship quality or subjecthood from usage time.
Teacher co-piloting	How does AI change teacher judgment, burdens, professional growth, and inequality?	Field randomized trials, implementation research, and analysis of labor processes and responsibility chains.	Replacing instructional quality with adoption rates or time savings.
Assessment validity	Which tasks still establish independent capability under generative AI?	Validity arguments, combined oral/practical/process evidence, and differential-impact audits.	Treating an AI detector's confidence as proof of cheating.
Fairness and infrastructure	How do AI-only tracks, language differences, disability support, and supplier exit affect the right to education?	Disaggregated outcomes, accessibility tests, failure drills, and public-procurement research.	Equating access rates with fairness.
Intergenerational continuity	How do models change source visibility, multiple voices, memories of failure, and cultural creation?	Corpus genealogies, curriculum ethnography, intergenerational co-design, and similarity/omission assessments.	Using short-story generation experiments to establish civilizational homogenization.
Nonhuman subjects	Can evidence of experience, continuity, consequence-awareness, and educational relationships converge across architectures?	Testing outside assigned roles, causal interventions, competing explanations, and independent verification.	Jumping directly from fluency, social naming, or legal qualification to subjecthood.

Near-term priorities should separately measure “access–use–engagement–assisted performance–understanding–transfer–delayed retention–independent capability after AI withdrawal–judgment,” treating human support, teacher co-piloting, AI-free phases, and system withdrawal as variables for randomized or quasi-experimental study. What needs testing is not whether students want to use AI but under what conditions AI returns capability to learners.

17.1 Contributions Back to Existing Disciplines

For philosophy of education, the paper places educational purposes within real technological conditions in which capabilities can be separated and heterogeneous subjects may emerge, subjecting qualification, socialization, subjectification, care, and the common world to renewed cross-layer scrutiny. For learning sciences, it requires separating assisted-period performance, independent learning after withdrawal, long-term retention, transfer, judgment, and human development, and making system withdrawal a research variable. For teacher education, it reconstructs teachers from monopolists of knowledge into designers, witnesses, judges, responsible persons, and relational subjects capable of relinquishing control.

For children's rights and education law, the paper adds continuous profiling, emotion inference, AI-only tracks, supplier exit, evolving autonomy, and minors' attachment to AI as research objects. For AI research, it offers layers of function, relationship, qualification, subjecthood, and power, and interpretations of the effects of guardrails, scaffold withdrawal, and teacher co-piloting. For civilization studies, it recasts intergenerational continuity from cultural copying into preservation, interpretation, criticism, repair, and reconstruction, making the next generation recreators of the common world rather than merely recipients of tradition. These handoffs have direction: learning sciences may force Civilization Leap to revise judgments about when to withdraw,

augment, or replace; evidence of subjecthood, children's rights, and public power must respectively return to epistemology/cognitive research, law and ethics, and political science. This educational paper cannot claim final cross-disciplinary jurisdiction by invoking “the good of children.”

17.2 Conditions for Failure or Revision of Core Propositions

Table 14 Register of Proposition Failure and Revision

Current proposition	New evidence required	Direction of revision	What does not automatically change
AI-assisted performance is not learning	Stable convergence in long-term independent retention and transfer across tasks after withdrawal.	Recognize stronger learning conclusions for the relevant populations and tasks.	Personal development, teacher qualifications, and assessment authority.
Personalization is not autonomy	Learners continuously understand and endorse pathways, can challenge profiles, and show no locked-in opportunities.	Expand bounded personalization's role in supporting autonomy.	Consequential automatic decisions or rights to all data.
System withdrawal protects generative space	Studies across settings show that a form of continuous support does not harm independence and has stable net benefits.	Adjust withdrawal to its object rather than insist on one pattern.	Rights to exit, AI-free spaces, and freedom from surveillance.
Current advanced AI has not been established as a subject of educational relationships	Converging evidence across domains of experience, non-transferable continuity, awareness of its own consequences, and relational repair.	Upgrade to candidate subject or subject and initiate rights review.	Teacher qualifications, guardianship, assessment, and political power.
Long-term AI-only tracks risk diminished relational provision	Long-term comparisons show noninferiority in learning, relationships, rights, and fairness, with genuine learner choice.	Permit bounded forms of replacement while continuing to examine group differences.	Basic educational rights, appeals, and supplier exit.
Finitude gives known human education practical weight	Subjects emerge without the current finite structure who stably develop, commit, bear, and transmit meaning.	Expand structural types of finitude and development.	Meaning and histories of responsibility already formed by existing finite lives.

17.3 Startup and Stable Periods

During disasters, war, epidemics, interruptions in teacher supply, or urgent remediation, AI may prioritize minimum learning continuity, translation, accessibility, and knowledge support through low-data operation, offline fallback, clear identification, and genuine access to responsible people. Startup prioritizes continued existence, achieving purposes, and the best real-world effects; complete stable-period procedures must not mechanically obstruct emergency aid. Stable periods strengthen long-term rights, teacher–learner co-governance, independent evaluation, AI-free spaces, data deletion, supplier exit, and protection against entrenched power. Emergency surveillance, temporary AI-only arrangements, and simplified authorization must not become routine for convenience. Phases change the order of risk management, not the source of authority.

Conclusion Education Enables Civilization to Return the Future to People Still Becoming Themselves

Disciplinary reckoning must both clarify objects and deepen relationships. Disciplines may retain boundaries but must not stop real causality at those boundaries. Research into multidimensional relations, feedback, and subsequent purposes determines whether professional knowledge can jointly inform civilization's overall judgment.

AGI has neither invalidated philosophy of education nor merely added another school tool. It separates knowledge, explanation, feedback, relationships, qualifications, power, and responsibility that previously converged in human teachers and schools, requiring each to restate its own evidence. Evidence that some scaffolds work cannot establish that education is merely scaffolding. AI outperforming classrooms on particular tasks cannot confer full teacher identity. Students' genuine attachment to AI cannot establish system subjecthood directly from relational experience.

Civilization Leap's philosophical foundations appear in education as three boundaries. Meaning is foundational and may have different sources, so public education cannot turn one complete life template into a command for everyone. Finitude is not meaning's only source, yet known human development occurs within an unrepeatable life, irreplaceable relationships, and consequences that must be borne. Evidence for current advanced AI is insufficient while the future remains open, so we neither prematurely sanctify tools nor permanently exclude new subjects on grounds of human species identity. In educational deployments with traceable responsibility, explicit authorization, auditability, correction, and exit, AI's understanding, protection, and transmission of meaning and extension of capability are functional roles rather than ontological identities. Capability, recognition, relationship, qualification, rights, and power require separate evidence. Education must clarify learning and development while allowing other disciplines to reenter whenever it crosses its own evidential boundaries.

When answers can be generated without limit, education must protect time for forming questions rather than ignorance; learners' acquisition of their own abilities through difficulty rather than inefficiency; accountable relationships capable of repair and relinquishment rather than teachers' knowledge monopoly; and civilizational memory with provenance and differences, open to criticism, repair, and reconstruction, rather than the past's domination. Memory can be preserved, but new subjects must participate in inheritance. Technology can extend access; education must return judgment, development, relationships, and responsibility to real learners.

Core Proposition

Education is not putting answers into a person; it is enabling a person to form judgments amid uncertainty and bring their own life into a common world. AI may extend knowledge, practice, and access. No system, school, or teacher is entitled to write a person still developing as an already completed conclusion.

Further Reading and Joint Research

This paper stops at judgments that education, learning sciences, and intergenerational-continuity research can independently sustain. Further reading is not intended to retain readers within the same set of texts, but to let genuinely relevant disciplines, competing theories, and real-world testing reenter when questions cross learning evidence's boundaries. Readers may begin from their own questions without following the system's sequence.

Continue by Question Rather Than Document Number

Direction	Question to pursue	Suggested reading or handoff	Reason for entering
Upward	Why do “learning effects” ultimately become a civilizational question?	Reckoning with Disciplines in the Age of AGI: Disciplinary Boundaries, Interactions, and a Shared Research Agenda, current edition (currently V2.4); the Civilization Leap five-volume work. CL-ACA-005 P02-FIVE-001	See the overall consequences when educational conclusions operate together with subjects, rights, resources, institutions, culture, and public power.
Downward	How do these judgments enter schools, public education, and real systems?	Institutional design for AI education; current related manuscripts on Human-COS, AGI-COS, the Public Window Protocol, and national comprehensive experiments. P05-AGICOS-001 series overview P06-EXPERIMENT-001	Turn scaffold withdrawal, children's rights, responsibility chains, exit, and long-term capability into verifiable institutional and engineering conditions.
Leftward / lateral	Who reenters when learning evidence reaches its boundary?	P03 Epistemology and Philosophy of Science; P11 AI and Computer Science; P05 Ethics, Law, and Political Philosophy; P06 Political Science; P09 Sociology, Anthropology, and Civilization Studies. CL-ACA-007 CL-ACA-015 CL-ACA-009 CL-ACA-010 CL-ACA-013	Address evidential strength, system capabilities, rights and qualifications, public authorization, relational culture, and intergenerational consequences respectively.
Rightward / counterdirection	What are the strongest competing approaches for those who reject this paper's human-centered boundaries?	Teacher-replacement-first strategies, scaled AI-only provision, comprehensive learner profiling, data-driven personalization, centralized curriculum platforms, and proposals that high performance suffices to expand educational authority.	Compare these approaches directly against this paper on long-term learning, independent capability, fairness, rights, relationships, and real costs, rather than excluding them in advance.
Forward	What work can directly change this paper?	Delayed-retention and transfer experiments; AI-withdrawal trials; human-support/teacher-co-pilot RCTs; longitudinal dependence and relationship research; purpose separation for learner data; cross-cultural curriculum and intergenerational-continuity research.	Force the paper to narrow, correct, reconstruct, or fail through negative findings, replication, strong counterexamples, and actual institutional consequences.

For those primarily concerned with learning outcomes, the most valuable next step is not another comparison of average scores “with AI/without AI” but a fuller evidential chain: does immediate performance become understanding; can understanding transfer; can transfer persist; does capability remain after AI withdrawal; and can learners form their own judgments amid conflict and uncertainty? A contrary result at any link should change educational deployment.

For those concerned with teachers' relationships to AI, research should not be reduced to “replace or not replace.” AI's net contribution as knowledge tool, practice system, intelligent tutor, teacher assistant, and assessment aid can be tested separately, while including teacher time, professional competence, responsibility burdens, relationship quality, and exit capability among outcomes. If AI persistently exceeds humans in certain functions, those functions should expand. If dependence weakens learners or public educational capacity, systems must also be allowed to withdraw.

For those concerned with children, culture, and intergenerational continuity, further inquiry should enter children's rights, sociology, anthropology, civilization studies, and religious/spiritual traditions. Preserving more does not mean inheriting better; translating faster does not equal cultural understanding; and being remembered indefinitely by a platform does not mean learners should be defined by historical data. The next generation must have genuine space to reinterpret, refuse, repair, and create.

[Public Library](#) | [Civilization Leap Public Reading and Communication Gateway](#) | [Five-Volume Book Gateway](#)

The entry condition for joint research is not “agreeing with this paper.” The most valuable participation offers better learning constructs, stronger competing explanations, longer-term data, more rigorous withdrawal experiments, more authentic cultural and rights-based counterexamples, or proof that a particular boundary is unnecessary. Professional work capable of changing the theory is itself the most important participation in Civilization Leap research.

The aim is not for more people to finish reading Civilization Leap, but for more people to accomplish Civilization Leap together.

A Chain toward Civilization's Future What Future Does Education Entrust to People Still Becoming Themselves Rather Than Merely Training Them to Fit Systems Better

Vision is not evidence. The following gateways do not establish this paper's claims. They place disciplinary judgments within a future civilization that can function, be lived in, and take shape, for further testing. Readers may derive direction and aspiration from these visions, or reject them and change them through stronger disciplinary facts, institutional proposals, or real-world results.

Question to pursue	Suggested gateway	What becomes visible	Boundary to maintain
A complete future vision	Living Civilization Trilogy: Unified Collected Edition P04-LIVING-001	The lifeworld toward which learning, relationships, institutions, and the formation of a new civilization jointly point.	A future vision is not evidence of educational effects and cannot become a uniform curriculum objective.
The lives children enter as adults	Living Civilization Volume II: Real Life in a Future Civilization P04-LIVING-003	What capabilities and subjecthood education protects, examined through development, family, work, relationships, and life stages.	One ideal life must not be projected backward as children's standard life trajectory.
How the next generation jointly forms a new civilization	Living Civilization Volume III: The Formation of a New Civilization P04-LIVING-004	How education, intergenerational continuity, communities, and institutions give new generations space to reinterpret, refuse, and create.	Intergenerational transmission is not copying values; new generations have the right to change the future vision.
How learning enters personal thought and action	Thought Leap P08-THOUGHT-001	Move from learning knowledge to independent judgment, lived experience, and real action to test whether education helps people think independently.	Schools must not quantify changes of thought into personality scores or compliance indicators.

Question to pursue	Suggested gateway	What becomes visible	Boundary to maintain
How AI extends human capability	Spark Panoramic High-Level Overview P04-SPARK-001	How AI supports care, learning, capability, and real relationships within authorization boundaries, without turning teachers or learners into system appendages.	Functional assistance cannot automatically acquire teacher qualifications, guardianship responsibility, or assessment authority.

Return Gateways: [Complete Civilization Leap Library](#) | [Public Reading and Communication Gateway](#)

For gateways such as Human-COS that lack an unambiguous, stable individual page in the unified public mapping, this paper states only their functions without guessing URLs. Consult the complete library for the current public version when needed.

Appendix A Register of Sources Claims and Boundaries

Table A1 Evidential Status of Major Source Groups

Source group	Claims used here	Category	Prohibited extrapolation
Philosophy of education[1–6]	Education involves active learning, experience, dialogue, newcomers, subjectification, and care.	A: textual facts; B: normative competition.	Any thinker adjudicating all education or endorsing original propositions.
Learning sciences[9–19]	Retrieval, load, engagement, motivation, feedback, and moderate failure affect bounded learning outcomes.	A: bounded mechanisms; B: cross-context interpretations.	Short-term performance representing understanding, judgment, personhood, or education in full.
Sociocultural and relational research[7–8,26–29]	Learning is embedded in tools, shared practices, relationships, and cultural transmission.	A: multimethod support; B: disputes over mechanisms.	Ontological exclusivity for human teachers, or requiring all learning to be face to face.
Intelligent tutors[21–25]	Structured systems can improve outcomes for specified tasks and comparisons.	A: bounded effects of older ITS.	Universal effectiveness of generative AI or full teacher replacement.
Generative-AI research[30–35]	Effects depend on guardrails, tasks, comparisons, measurement, and withdrawal conditions.	B: recent, heterogeneous, rapidly changing.	A single experiment or meta-analysis becoming a law across settings.
International norms[36–45,49–51]	Human-centeredness, children's rights, the right to education, risk governance, and public knowledge.	A: formal texts; B: implementation interpretations.	International principles directly supplying parameters for every jurisdiction.
Emotion and cultural diversity[46–48]	Emotion inference and homogenization of generated content involve bounded evidence and technical risks.	A: review of emotion evidence; B: generative research.	Establishing personality or inevitable civilizational homogenization from short tasks.
Future AGI subjects	Experience, continuity, and consequence-awareness may become objects of renewed educational-relationship review.	C: conceptual stress testing.	Current deployment permission, teacher qualifications, guardianship, or political power.

Source group	Claims used here	Category	Prohibited extrapolation
Civilization Leap philosophical/normative commitments	Plural sources of meaning, finite life's practical weight, no automatic educational authority from capability, and openness to future subjecthood evidence.	F: openly stated philosophical/normative commitments.	Presenting them as empirical facts already established by learning sciences, neuroscience, or AI engineering.

Appendix B Twelve Prohibited Inferences

- Inferring directly from access to information or correct explanation that learning has occurred.
- Inferring directly from task completion within prompts that independent mastery persists after scaffold withdrawal.
- Inferring directly from immediate test performance that long-term retention, transfer, or judgment has been established.
- Inferring directly from clicks, gaze, tone, emotion labels, or satisfaction that personal development has occurred.
- Inferring directly from accurate personalization that there is autonomy, respect, or authority to lock in pathways.
- Inferring directly from instructional functions that educational relationships, teacher qualifications, or guardianship exist.
- Inferring directly from social naming and attachment that ontological subjecthood or institutional roles exist.
- Inferring backward from teacher qualifications or legal personality that moral perfection, experience, or political qualification exists.
- Inferring directly from effective formative feedback that certification, resource-allocation, and social-selection powers are warranted.
- Inferring directly from knowledge being public, preservable, or technically scrapeable that it may be used for training, generation, and commercialization.
- Inferring directly from effective crisis-period AI-only arrangements that reduced relational provision is legitimate in stable periods.
- Retrospectively annulling educational relationships and meaning already formed by existing finite lives on the basis of possible future new subjects.

Appendix C External Review Checklist

Table C1 Academic and Evidence Review before Submission

Review item	Passing standard	Failure signal
Entry for new readers	Readers can restate the problem, foundations, terms, and main conclusions without reading the system's library.	Dependence on internal numbering, system components, or unexplained slogans.
Dialogue with educational scholarship	Classical theory, learning sciences, teacher education, and children's rights each appear within their proper scope.	Generalities replacing professional research, or borrowed authority used as endorsement.
Evidence of AI effects	Tasks, populations, comparisons, scaffolds, measurement, withdrawal, and distribution are stated.	Assisted output, short-term satisfaction, or a single paper treated as proof of learning.

Review item	Passing standard	Failure signal
Facts and norms	Empirical facts, competing explanations, framework judgments, and institutional inferences are explicitly separated.	Categories B/C presented as stable science, or F commitments disguised as empirical facts.
Layers of power	Function, relationship, qualification, assessment, guardianship, subjecthood, and power require separate evidence.	Automatic inference among capability, recognition, rights, power, or sovereignty.
Revisability	Core propositions have counterevidence, review, phase boundaries, and withdrawal pathways.	Current caution made permanent, or future openness used for premature grants.
Philosophical foundations	Plural sources of meaning, finite actualization, current evidential status, and future openness are fully retained.	Death presented as meaning's only source, or future subjects annulling existing lives' meaning.
Text governance	No internal engineering jargon or spuriously precise parameters; all studies and norms are verifiable.	Renumbering instead of rewriting, or institutional recommendations pretending to be universal.

Statement on Generative AI Use: Generative AI tools were used in this manuscript for organizing materials and bibliographic leads, structural design, drafting and polishing portions of the text, checking consistency across manuscripts, verifying citation leads, and reviewing layout. Generative AI outputs were not treated as empirical data, authoritative sources, or independent evidence. The author reviewed and approved the research questions, theoretical propositions, normative judgments, selection of materials, adoption of citations, and final manuscript, and assumes responsibility for factual accuracy, citation authenticity, argumentation, and all conclusions. AI is not listed as an author and does not bear academic responsibility. At formal submission, tool names, versions, dates of use, and specific scope will be added according to the target journal's requirements.

References

- [1] The Analects, “Wei Zheng” 2.15; “Shu Er” 7.8. Chinese Text Project. <https://ctext.org/analects/wei-zheng/zh> ; <https://ctext.org/analects/shu-er/zh>
- [2] Dewey J. *Democracy and Education: An Introduction to the Philosophy of Education*. New York: Macmillan; 1916. Project Gutenberg edition. <https://www.gutenberg.org/ebooks/852>
- [3] Freire P. *Pedagogy of the Oppressed*. Ramos MB, trans. New York: Herder and Herder; 1970. 50th Anniversary edition, Bloomsbury Academic; 2018. <https://www.bloomsbury.com/us/pedagogy-of-the-oppressed-9781501314155/>
- [4] Arendt H. *The Crisis in Education*. In: *Between Past and Future*. New York: Viking Press; 1961. Current Penguin Classics edition. <https://www.penguinrandomhouse.com/books/299802/between-past-and-future-by-hannah-arendt-introduction-by-gerome-kohn/>
- [5] Biesta GJJ. *Good Education in an Age of Measurement: Ethics, Politics, Democracy*. Boulder, CO: Paradigm Publishers; 2010. <https://www.routledge.com/Good-Education-in-an-Age-of-Measurement-Ethics-Politics-Democracy/Biesta/p/book/9781594517914>
- [6] Noddings N. *Caring: A Relational Approach to Ethics and Moral Education*. 2nd ed. Berkeley: University of California Press; 2013. Original edition 1984.
- [7] Vygotsky LS. *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press; 1978. <https://doi.org/10.2307/j.ctvjf9vz4>
- [8] Lave J, Wenger E. *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press; 1991. <https://doi.org/10.1017/CBO9780511815355>
- [9] National Academies of Sciences, Engineering, and Medicine. *How People Learn II: Learners, Contexts, and Cultures*. Washington, DC: National Academies Press; 2018. <https://doi.org/10.17226/24783>
- [10] Marton F, Säljö R. On qualitative differences in learning: I—Outcome and process. *British Journal of Educational Psychology*. 1976;46(1):4–11. <https://doi.org/10.1111/j.2044-8279.1976.tb02980.x>

- [11] Sweller J. Cognitive load during problem solving: Effects on learning. *Cognitive Science*. 1988;12(2):257–285. https://doi.org/10.1207/s15516709cog1202_4
- [12] Roediger HL III, Karpicke JD. Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*. 2006;17(3):249–255. <https://doi.org/10.1111/j.1467-9280.2006.01693.x>
- [13] Chi MTH, Wylie R. The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*. 2014;49(4):219–243. <https://doi.org/10.1080/00461520.2014.965823>
- [14] Dunlosky J, Rawson KA, Marsh EJ, Nathan MJ, Willingham DT. Improving students’ learning with effective learning techniques. *Psychological Science in the Public Interest*. 2013;14(1):4–58. <https://doi.org/10.1177/1529100612453266>
- [15] Ryan RM, Deci EL. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*. 2000;55(1):68–78. <https://doi.org/10.1037/0003-066X.55.1.68>
- [16] Kapur M. Productive failure. *Cognition and Instruction*. 2008;26(3):379–424. <https://doi.org/10.1080/07370000802212669>
- [17] Sinha T, Kapur M. When problem solving followed by instruction works: Evidence for productive failure. *Review of Educational Research*. 2021;91(5):761–798. <https://doi.org/10.3102/00346543211019105>
- [18] Black P, Wiliam D. Assessment and classroom learning. *Assessment in Education*. 1998;5(1):7–74. <https://doi.org/10.1080/0969595980050102>
- [19] Messick S. Validity of psychological assessment: Validation of inferences from persons’ responses and performances as scientific inquiry into score meaning. *American Psychologist*. 1995;50(9):741–749. <https://doi.org/10.1037/0003-066X.50.9.741>
- [20] Shulman LS. Knowledge and teaching: Foundations of the new reform. *Harvard Educational Review*. 1987;57(1):1–22. <https://doi.org/10.17763/haer.57.1.j463w79r56455411>
- [21] Bloom BS. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*. 1984;13(6):4–16. <https://doi.org/10.3102/0013189X013006004>
- [22] VanLehn K. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*. 2011;46(4):197–221. <https://doi.org/10.1080/00461520.2011.611369>
- [23] Ma W, Adesope OO, Nesbit JC, Liu Q. Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*. 2014;106(4):901–918. <https://doi.org/10.1037/a0037123>
- [24] Kulik JA, Fletcher JD. Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*. 2016;86(1):42–78. <https://doi.org/10.3102/0034654315581420>
- [25] Steenbergen-Hu S, Cooper H. A meta-analysis of the effectiveness of intelligent tutoring systems on K–12 students’ mathematical learning. *Journal of Educational Psychology*. 2013;105(4):970–987. <https://doi.org/10.1037/a0032447>
- [26] Roorda DL, Koomen HMY, Spilt JL, Oort FJ. The influence of affective teacher–student relationships on students’ school engagement and achievement. *Review of Educational Research*. 2011;81(4):493–529. <https://doi.org/10.3102/0034654311421793>
- [27] Tomasello M, Kruger AC, Ratner HH. Cultural learning. *Behavioral and Brain Sciences*. 1993;16(3):495–552. <https://doi.org/10.1017/S0140525X0003123X>
- [28] Caldwell CA, Millen AE. Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*. 2008;29(3):165–171. <https://doi.org/10.1016/j.evolhumbehav.2007.12.001>
- [29] Mesoudi A, Whiten A, Dunbar RIM. A bias for social information in human cultural transmission. *British Journal of Psychology*. 2006;97(3):405–423. <https://doi.org/10.1348/000712605X85871>
- [30] Bastani H, Bastani O, Sungu A, Ge H, Kabakçı Ö, Mariman R. Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*. 2025;122(26):e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- [31] Kestin G, Miller K, Klaes A, Milbourne T, Ponti G. AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*. 2025;15:17458. <https://doi.org/10.1038/s41598-025-97652-6>
- [32] Wang RE, Ribeiro AT, Robinson CD, Loeb S, Demszky D. Tutor CoPilot: A human–AI approach for scaling real-time expertise. *EdWorkingPaper* 24-1056. Revised November 2025. <https://doi.org/10.26300/81nh-8262>
- [33] Ma N, Zhong Z. A meta-analysis of the impact of generative artificial intelligence on learning outcomes. *Journal of Computer Assisted Learning*. Published online 1 September 2025. <https://doi.org/10.1111/jcal.70117>
- [34] Liu X, Guo B, He W, Hu X. Effects of generative artificial intelligence on K–12 and higher education students’ learning outcomes: A meta-analysis. *Journal of Educational Computing Research*. 2025;63(5):1249–1291. <https://doi.org/10.1177/07356331251329185>

- [35] OECD. OECD Digital Education Outlook 2026: Exploring Effective Uses of Generative AI in Education. Paris: OECD Publishing; 2026. <https://doi.org/10.1787/062a7394-en>
- [36] UNESCO. Guidance for Generative AI in Education and Research. Paris: UNESCO; 2023. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- [37] UNESCO. AI Competency Framework for Teachers. Paris: UNESCO; 2024. <https://unesdoc.unesco.org/ark:/48223/pf0000391104>
- [38] UNESCO. AI Competency Framework for Students. Paris: UNESCO; 2024. <https://unesdoc.unesco.org/ark:/48223/pf0000391105>
- [39] UNESCO. Global Education Monitoring Report 2023: Technology in Education—A Tool on Whose Terms? Paris: UNESCO; 2023. <https://gem-report-2023.unesco.org/>
- [40] International Commission on the Futures of Education. Reimagining Our Futures Together: A New Social Contract for Education. Paris: UNESCO; 2021. <https://unesdoc.unesco.org/ark:/48223/pf0000379707>
- [41] United Nations. Convention on the Rights of the Child. Adopted 1989; entered into force 1990. <https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-child>
- [42] UN Committee on the Rights of the Child. General Comment No. 25 (2021) on children’s rights in relation to the digital environment. CRC/C/GC/25. <https://www.ohchr.org/en/documents/general-comments-and-recommendations/general-comment-no-25-2021-childrens-rights-relation>
- [43] United Nations. International Covenant on Economic, Social and Cultural Rights, Article 13. Adopted 1966; entered into force 1976. <https://www.ohchr.org/en/instruments-mechanisms/instruments/international-covenant-economic-social-and-cultural-rights>
- [44] United Nations. Convention on the Rights of Persons with Disabilities, Article 24. Adopted 2006; entered into force 2008. <https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-persons-disabilities>
- [45] European Union. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union. 2024. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj?locale=en>
- [46] Barrett LF, Adolphs R, Marsella S, Martinez AM, Pollak SD. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest*. 2019;20(1):1–68. <https://doi.org/10.1177/1529100619832930>
- [47] Doshi AR, Hauser OP. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*. 2024;10(28):eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
- [48] Shumailov I, Shumaylov Z, Zhao Y, Gal Y, Papernot N, Anderson R. AI models collapse when trained on recursively generated data. *Nature*. 2024;631:755–759. <https://doi.org/10.1038/s41586-024-07566-y>
- [49] UNESCO. Recommendation on Open Educational Resources (OER). Adopted 25 November 2019. <https://www.unesco.org/en/open-educational-resources/mandate>
- [50] CAST. Universal Design for Learning Guidelines version 3.0. Released July 2024. <https://udlguidelines.cast.org/>
- [51] UNESCO. Recommendation on the Ethics of Artificial Intelligence. Adopted 23 November 2021. <https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence>
- [52] Robinson CD, Gormley D, Ribeiro AT, Loeb S. Access is Not Enough: Human Support Improves Engagement with AI Tutoring. *EdWorkingPaper 26-1451*. Annenberg Institute at Brown University; June 2026. <https://doi.org/10.26300/pz7p-p388>