

The Well-Formed Mind *in the AI Age*

What education must preserve, and how AI can help



Illustration: books, reasoning and human relationships surround the learner; AI supports their development.

The decisive test comes after assistance ends: what can the learner still know, explain, verify and do?

Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).

THE CENTRAL QUESTION

What remains when the screen is closed?

A polished answer is easy to admire. The harder educational question comes when the screen is closed: can the learner explain the answer, examine its sources, apply the idea elsewhere and recognize when it is wrong? That is the test at the center of this investigation. Artificial intelligence changes the means of producing schoolwork; it does not remove the need to establish what the person has learned.

The stakes extend beyond academic integrity. Education prepares people to judge claims, govern their attention, accept responsibility and work with others. A tool is educationally valuable when it helps develop those capacities. Faster completion can be useful, but the submitted product and the person's growing capability are different outcomes.

The evidence reviewed for this investigation supports bounded uses of artificial intelligence that provide questions, hints, feedback, source comparison, accessibility, or assistance to a responsible human educator. It does not support treating unrestricted answer generation as education.

A strong student product cannot prove learning when the tool may have supplied the reasoning, language, or solution. The decisive test comes later: can the student understand, explain, verify, and apply the knowledge without the tool?

The investigation's argument is that AI makes an older weakness visible: institutions may specify subjects, credits and approved technology without agreeing on the independent human abilities all of them should form. The remedy begins with educational purpose, then asks how teaching, assessment and technical governance can serve it together.

The reading route

Human purpose and the institutional gap

What the evidence shows

A protected-learning model

Authorship, the whole person and California

Current determination and research record

FOUNDATIONS

What education should form

Socrates begins with examination. Education should not merely give a student information or confidence; it should cultivate the discipline to question claims, test reasons, recognize the limits of personal knowledge, and distinguish genuine understanding from reputation, fluency, or unsupported certainty.

Plato extends the question from what a person knows to how the whole person is ordered. In the *Republic*, education is a turning toward what is real and good. Reason, desire, emotion, attention, courage, and conduct must be formed together. A curriculum that transmits facts while leaving judgment, appetite, responsibility, and character untouched is therefore incomplete.

This investigation does not treat either philosopher as supplying a complete modern program. It uses their enduring questions as the starting test: does education form people who can seek truth, govern themselves, reason independently, accept correction, and act responsibly? Artificial intelligence makes that question more urgent because a tool can produce impressive work without establishing that

the learner possesses the knowledge, judgment, or capability the work appears to demonstrate.

A well-formed mind is not a storehouse of facts, a high test score, a political identity, or technical fluency alone. It combines knowledge with the habits needed to examine evidence, govern attention, recognize uncertainty, make moral distinctions, communicate honestly, and participate responsibly with other people. The standard draws from classical philosophy, liberal education, Great Books, scientific and mathematical reasoning, developmental learning, and modern work on the integration of knowledge. No single thinker supplies the whole definition.

Daniel Huttenlocher's questions about human-AI collaboration and changes in human behavior offer a useful frame for this investigation: verify what the learner can still explain, judge and do independently after AI assistance ends. The questions do not establish educational benefit; learning outcomes require direct studies and observed performance.

MIT : Generative AI Impact Consortium
(Huttenlocher's research questions).



Illustration: a thinker pauses over an open book, reflecting the examination of ideas at the heart of education.

FROM IDEAL TO OBSERVABLE PRACTICE

A mind that can act on its own

Foundational knowledge

Recall and use essential facts, concepts, language, and procedures without constant external retrieval

Reasoning

Follow premises, detect contradictions, solve unfamiliar problems, and explain each major step

Truth verification

Trace claims to sources, distinguish fact from inference, and revise a conclusion when evidence changes

Writing and speech

Express an intelligible position in the student's own voice and answer questions about it

Intellectual humility

Recognize uncertainty, limits, error, and the strongest reasonable objection

Moral judgment

Consider purpose, responsibility, consequences, human dignity, and the common good

Attention and perseverance

Sustain effort through productive difficulty rather than escaping every demanding mental act

Self government

Choose when to use a tool, disclose assistance, and refuse uses that weaken learning or responsibility

Human relationship

Listen, disagree constructively, accept correction, cooperate, and remain accountable to other people

Independent performance

Demonstrate the intended knowledge or skill after assistance is removed

The standard changes with age. Young children need language, memory, direct experience, imitation, play, stories, number sense, and trusted adults. Older students progressively add argument, abstraction, rhetoric, scientific method, moral responsibility, and independent synthesis. Adults should be able to integrate specialized knowledge with judgment about purpose and consequence. Artificial intelligence should not enter every stage in the same way.

THE PROBLEM AI EXPOSES

Many authorities, one unfinished task

Modern education contains excellent teachers, demanding courses, serious institutions, and successful students. Its governing structure, however, rarely unites curriculum, assessment, technology, family responsibility, and institutional accountability around one explicit objective: forming people who can know, remember, reason, judge, and act well. Artificial intelligence makes that weakness harder to ignore because it can support the student's work or quietly perform the work in the student's place.

The central problem is not the absence of educational authority. Public education has extensive authority distributed among legislatures, state agencies, governing systems, districts, boards, administrators, faculty bodies, teachers, vendors, families, and students. The problem is that this structure does not consistently coalesce around an observable account of the mind and character education should form. Distributed authority can protect local judgment. It can also disperse responsibility until no actor owns the full result.


Who shapes the learner?

Parents

Values, daily habits and the conditions for learning.

Teachers

Tasks, feedback and opportunities to explain.



School leaders

Curriculum, assessment, access and tool rules.

Technology providers

Design, safeguards and the terms of use.

Parents, teachers, school leaders and technology providers all influence what the student learns. Their decisions need to work together, with a clear focus on what the student can understand and do independently.

WHERE THE HANDOFFS WEAKEN

One purpose, many decisions

- Curriculum can become a collection of subjects and requirements without a shared account of how they develop memory, reasoning, judgment, and responsibility.
- Assessment often rewards the completed product more clearly than the student's unaided understanding or capacity to defend it.
- Teacher specialization can provide depth while leaving educators without a common formation standard or reliable evidence about unfamiliar methods.
- Parents may remain legally responsible while lacking time, confidence, access, or effective influence over curriculum and technology decisions.
- Administrative, political, philanthropic, commercial, and ideological incentives can shape education without being tested against the same human outcome.
- Technology procurement can move faster than course design, privacy review, faculty governance, developmental limits, and outcome measurement.

Alternatives and their limits

Homeschooling, classical schools, independent schools, charter models, religious education, project-based learning, and Great Books programs can restore clearer responsibility or educational purpose. They do not automatically solve the problem. Each model still requires capable adults, sound curriculum, intellectual honesty, social support, accountability, accessibility, and evidence that students can perform independently.

The wider investigation examines the philosophies and movements that changed educational continuity, including contested modern figures and political methods. Those comparisons should clarify influence and consequence without treating disagreement as proof of misconduct or reducing education to a partisan contest.



Illustration: adults discuss learning materials together, connecting decisions about teaching, resources and daily support with the learner's needs.

WHAT THE STUDIES SHOW

Better work is not yet better learning

The research does not yield a single verdict on “AI in education.” It distinguishes modes of help, particular tasks and the moment at which learning is measured. An assistant that supplies an answer, a tutor that releases hints and a system that supports a human teacher are different interventions. Their findings should not be merged into a product endorsement.

When assistance is removed

A randomized study of nearly 1,000 high-school mathematics students compared no artificial intelligence, a standard GPT-4 interface, and a guarded tutoring interface. Practice performance increased with both AI conditions. When the tool was removed, students who had used the standard interface scored 17 percent below the group that never received AI. The guarded tutor largely avoided the loss. The result shows why assisted grades can conceal weaker independent learning and why the exact design of the tool matters. [16]

Study boundary: the 17% difference is relative, not percentage points. The guarded tutor’s later unaided result was statistically indistinguishable from control. This was one high-school setting.

Purpose-built tutoring

A peer-reviewed crossover experiment in a Harvard college physics course found that students using a


custom AI tutor learned more in less time than students receiving an in-class active-learning lesson on the same material. The tutor was built around established teaching practices, so the finding does not establish that an ordinary chatbot would produce the same result. [17]

Study boundary: the experiment covered two physics lessons and immediate post-tests, rather than long-term retention.

Supporting the human tutor

Tutor CoPilot tested a different relationship. Artificial intelligence supplied suggestions to human tutors rather than replacing them. Students were 4 percentage points more likely to master lesson topics, with larger gains among students assigned to lower-rated tutors. The study supports AI-plus-human instruction as a serious model, although its working-paper status and specific setting limit generalization. [18]

From assistance to understanding



Harvard’s two-lesson study found higher immediate test scores with a custom AI tutor [17]. The diagram shows our proposed teaching sequence: explain, practice and apply independently. The study did not test this full sequence or long-term retention.

THE LEARNING MECHANISM

Protect the mental act being learned

Memory still requires retrieval

Cognitive offloading can free attention for higher-order work, as notes, calculators, and search tools have long done. The risk arises when the external tool performs the retrieval, representation, or reasoning that the student is meant to learn. Research on transactive memory shows that expected access changes what people remember, while controlled research on retrieval practice shows that producing information from memory improves delayed recall. [19] [20]

What the evidence cannot yet establish

The strongest causal evidence concerns bounded tasks in mathematics, physics, programming,

writing, and instructional support. Long-term evidence about attention, independent judgment, relationships, moral development, anxiety, belonging, and family life remains thin. Satisfaction, convenience, engagement, and polished work should not be reported as proof of mental fitness or durable education. [21]

From task performance to capability

The potential benefit is more useful practice, feedback and access to explanation while the learner continues doing the work that builds understanding. The safeguard is to check that capability after assistance ends. This is a proposed educational design principle, not a guarantee that every learner or tool will benefit.

A stronger artifact

Clearer language, a completed solution or a faster finished task.

A stronger learner

Understanding that the person can explain, retain, verify and apply after answer-generating help ends.

Clearer work is one outcome; stronger understanding is another. The learner must still be able to explain, remember, check and apply the ideas after help ends.

HOW IT COULD WORK

Use the tool to strengthen the learner

The educational mode matters more than the brand. A general assistant may act as a tutor, critic, source guide, answer generator, companion, or substitute author. Suitability depends on the role it performs, the learner's age and preparation, the information it receives, and the independent performance required afterward.

01

Attempt

Begin with an independent first attempt where it serves the learning objective.

02

Receive bounded help

Use a question, hint, source comparison or critique. Keep the learner responsible for the solution.

03

Explain and verify

Describe the reasoning, check sources and correct errors. Preserve a proportionate account of material help.

04

Return without the tool

Use a later unaided explanation or new problem to check retention and transfer, with fair accommodations.

Begin with the learner's own attempt, add bounded help, then explain, verify and try again independently. This proposed teaching sequence keeps the learner doing the thinking and preserves necessary accommodations and access support. Read the verification standard.

MODE BEFORE BRAND

Choose the role. Protect the learning.

1. Human teaching and independent effort

The educator, parent, text, task, and learner remain primary

2. Source-bounded research

The learner can inspect the defined materials and trace every material claim

3. Guided tutoring

The system asks questions, releases hints gradually, and keeps the learner responsible for the solution

4. Critique, comparison and revision

The student attempts the work first and remains responsible for verification and substantive revision

5. Teacher support

The tool returns time or insight to the educator without making consequential human decisions

Uses that need strict limits

Limit AI that does the thinking an assignment is meant to teach, replaces human relationships, infers or tracks students' emotions, or makes high-stakes grading decisions. Exceptions need an educational reason, safeguards and meaningful human oversight.

The institutional minimum

- Name the human purpose of the course and each major assessment.
- Identify the knowledge and skills students must retain without artificial intelligence.
- Protect first attempts, retrieval, direct reading, original composition, live reasoning, oral defense, and practical performance where they serve the learning objective.
- Require traceable assistance, source verification, and student responsibility for the final work.
- Use approved accounts, defined data rules, accessibility review, and an equivalent route when AI is not essential.
- Expand a use only when independent learning improves after assistance is removed and serious subgroup harms are not concealed by the average.

The practical test is demanding: after using the tool, is the student more able to understand, explain, verify, choose, and act without it? [16] [17] [18]

The test above is a proposed criterion. The cited studies examine bounded learning outcomes, not every form of judgment or human flourishing.

What is the task meant to show?



Learn or practice

Use permitted questions, hints or critique. Keep the learner reasoning and checking.

Demonstrate capability

Use an independent explanation or performance, preserving fair accommodations.

The task determines the role of AI: hints and feedback can support practice; a demonstration of capability calls for independent work with appropriate accommodations.

MAKING THE STANDARD PRACTICAL

How the learner makes capability visible

- Require an independent first attempt when the learning objective permits it.
- Preserve drafts, source notes, calculations, prompts, or a concise process account when AI materially assists the work.
- Use oral explanation, live problem solving, supervised writing, practical performance, or delayed assessment to confirm understanding.
- Ask students to identify and correct errors rather than merely accept or reject an AI answer.




Illustration: students question, listen and explain ideas around a seminar table. Their discussion makes reasoning visible to the teacher and to one another.

RESPONSIBILITY

A student's name must still mean something

Artificial intelligence can summarize sources, propose explanations, generate citations, imitate confidence, and produce fluent error. Responsible education must therefore teach students to distinguish a plausible statement from a supported one. Verification requires opening the source, checking whether it says what the tool claims, identifying missing context, comparing strong contrary evidence, and accepting uncertainty when the record does not resolve the question.

A student's name on an assignment should still indicate responsibility for its claims, reasoning, sources, and language. Where assistance is permitted, the student should identify the tool and describe the material help it provided. Disclosure alone is insufficient when the assignment is intended to measure original composition, memory, reasoning, or judgment; those capacities also require a protected demonstration.

Personalized AI makes this problem sharper. The authorship strand of the investigation distinguishes systems that can already imitate aspects of a learner's language and knowledge from more capable future possibilities. It does not establish that a complete human proxy is inevitable. The educational response is to verify understanding through several proportionate forms of evidence rather than assume that a familiar writing style proves authorship. [Read the authorship record.](#)

A fair authorship review

- Judge misconduct from evidence about the work and applicable rule, not from an automated detector score alone.

Explain the claim. Trace the source.



Illustration: a student explains a claim and points to its source while a tutor listens. Discussing the reasoning and checking evidence make the student's understanding easier to assess.

Research has found that GPT-detection systems can disproportionately classify writing by non-native English writers as AI-generated. Detection may prompt review, but it should not substitute for fair procedure or evidence of authorship. [22]

The informed choice not to use AI

Competence includes knowing when artificial intelligence weakens privacy, attention, accuracy, authorship, or the purpose of an assignment. Students and educators should retain a meaningful non-AI route when the tool is not essential, subject to legitimate course requirements and accessibility needs.

Authorship is one part of a wider question: what happens to the learner's attention, judgment and relationships beyond the assignment?

ATTENTION, RELATIONSHIPS AND WELLNESS

The learner has a life beyond the assignment

This investigation uses mental fitness as a nonclinical educational term. It includes sustained attention, recall, tolerance of productive difficulty, explanation, judgment, self-regulation, responsibility, and healthy participation with other people. Each capacity requires its own evidence. Convenience and engagement are not substitutes.

Why COVID remains relevant

School closures and emergency remote education disrupted instruction, routine, friendship, family life, and access unevenly. Students who entered the disruption with stable relationships, sound habits, reliable instruction, and capable family support did not face the same conditions as students already experiencing isolation, weak routines, family dysregulation, limited friendships, disability barriers, or educational instability. These differences affect recovery and complicate later claims about technology. [23]

Mother of Divine Grace School, one of the institutions examined next, entered the pandemic with an established distance-learning and home-education structure. Its experience should not be treated as equivalent to emergency remote schooling adopted with little preparation. The comparison is useful precisely because delivery method alone cannot explain the quality of formation, family support, or social connection. [12]

Questions that remain open

- Does sustained AI use change attention, sleep, anxiety, confidence, belonging, loneliness, or willingness to struggle with difficult work?

Learning with and through others



Illustration: students build a model with a teacher, combining practical problem-solving, perseverance and cooperation. Learning also develops through shared effort and relationships.

- Do effects differ by age, disability, language, prior achievement, family support, or existing distress?
- Can AI provide remediation without becoming another form of avoidance or dependency?
- Does teacher or parent mediation change the result compared with unsupervised access?
- Do students maintain friendship, discussion, mentorship, and responsibility when more learning occurs through a machine?

Current evidence does not justify a general claim that artificial intelligence caused improvement or decline in these outcomes. The investigation treats them as required measures rather than predetermined conclusions.

The California review that follows applies these aims and limits to the published institutional record.

THE INSTITUTIONAL TEST

Purpose, policy and the evidence between them

Responsibility has to travel all the way to the classroom. A published mission can name the purpose, a policy can define permitted tools, and an assignment can make the rule concrete. The final question is whether the learner can demonstrate the intended capacity. Evidence at one stage cannot stand in for evidence at the next.

California's public systems show growing attention to human-centered use, access, governance, privacy, and integrity. Their strongest documents also reveal the unresolved problem: statewide or systemwide adoption can expand before institutions establish minimum protected human capabilities and comparable outcome measures.

The California Department of Education guidance and model policy are important reference points for K through 12 schools. The University of California Academic Senate report provides the strongest reviewed system-level warning that demanding intellectual work and faculty judgment cannot be replaced by technological enthusiasm. [1] [2] [5]

California Department of Education

Published strength. Guidance and a model policy address human learning, disclosure, original work, privacy, authorized tools, consent, family communication, bias, and well-being

Unresolved formation question. The model is nonmandatory and does not create one statewide minimum for unaided reading, writing, memory, calculation, oral defense, or judgment

California Community Colleges

Published strength. The system states that technology should amplify human capacity rather than replace it and connects AI with faculty leadership, ethics, equity, governance, and evaluation

Unresolved formation question. Pilots still need independent learning, dependency, subgroup, workload, and durable outcome measures

California State University

Published strength. Large-scale access can reduce unequal availability and support common training and security arrangements

Unresolved formation question. Procurement and workforce goals may move ahead of course purpose, faculty governance, protected assessment, privacy, and measured learning

University of California

Published strength. The Academic Senate report defends faculty authority, discipline-specific rules, transparency, opt-in choices, refusal, data protection, and consultation


Unresolved formation question. The system still needs clearer implementation evidence, academic-integrity rules, data ownership, and measures of reading, writing, creativity, and independent judgment

A GEOGRAPHIC READING GUIDE

The California case sample

The institutional review covers ten California cases. They represent public and private universities, a science and engineering institute, Catholic higher education, an independent secondary school, a public charter network, and a private Catholic K through 12 distance-learning and home-education program. The selection is comparative, not representative of every California school.

CALIFORNIA



SOUTHERN CALIFORNIA DETAIL



▲ K-12 ● Higher education

1 **Mother of Divine Grace**
Ojai admin. base; distance learning

2 **Harvard-Westlake**
Los Angeles; Upper School shown

3 **High Tech High**
San Diego County; Point Loma shown

4 **UC Berkeley**
Berkeley; main campus

5 **UCLA**
Los Angeles; Westwood

6 **Stanford University**
Stanford; main campus

7 **USC**
Los Angeles; University Park

8 **Caltech**
Pasadena; main campus

9 **Loyola Marymount University**
Los Angeles; Westchester

10 **Thomas Aquinas College**
Santa Paula; California campus

Approximate locations show where the ten cases are based, not the reach or effectiveness of their programs. Harvard-Westlake has two academic campuses; High Tech High is a multi-site network. Names link to official location records. Outline: U.S. Census Bureau, 2025.

THREE DIFFERENT MODELS

Protect foundational learning

Mother of Divine Grace School

Strongest public evidence. A parent-directed classical sequence explicitly develops memory, imagination, reasoning, communication, judgment, and responsibility

Evidence still needed. A complete public student-facing AI policy, approved tools, age and consent rules, privacy controls, implementation records, and measured outcomes

Harvard-Westlake School

Strongest public evidence. A demanding curriculum and visible ethical and dialogical learning provide a basis for original thought and constructive disagreement

Evidence still needed. Current schoolwide AI rules, approved accounts, parent notice, independent verification, accessibility, and comparable outcomes

High Tech High

Strongest public evidence. Authentic projects, public work, collaboration, and responsibility connect knowledge with practical creation

Evidence still needed. Detailed AI governance, privacy, age and consent, individual mastery checks, implementation consistency, and durable results

The three K through 12 institutions should be evaluated by developmental stage, parent authority, consent, family communication, student support, and protected foundational learning. This investigation does not authorize a student-facing AI condition in K through 12. [12] [14] [15]

Different models, a shared duty of care

Home education

Parent-directed sequence and daily formation.

Independent school

Demanding coursework and intellectual dialogue.

Project-based network

Authentic projects, public work and collaboration.

FOUNDATIONAL SAFEGUARDS

Developmental fit

Age, foundational learning and the learner’s needs.

Parent authority

Notice, consent and meaningful family communication.

Privacy and support

Approved tools, accessibility and student protection.

Independent evidence

Verify what the learner can understand and do.

The diagram brings three school models together around four checks: developmental fit, parent authority, privacy and support, and independent learning. Each setting needs evidence that these safeguards work.

SELECTED INSTITUTIONAL CASES

Strong practices, different evidence gaps

Colleges and universities should be evaluated by course purpose, faculty authority, disciplinary differences, independent scholarship, protected assessment, degree-level outcomes, and the student's capacity to assume adult responsibility. Their AI rules need not be identical, but students need a recognizable institutional floor.

UC Berkeley

Strongest public evidence. A mature AI hub, licensed tools, privacy guidance, assignment redesign, and course examples that protect independent reasoning

Evidence still needed. Campus-wide implementation, consistent protected capabilities, delayed learning, workload, subgroup, and wellness results

UCLA

Strongest public evidence. Student-use evidence, citation guidance, assignment-design resources, and institution-wide principles

Evidence still needed. A common independent-performance floor, approved-tool chain, policy implementation, and durable outcomes

Stanford University

Strongest public evidence. Human-centered AI literacy, student guidance, responsible-use controls, and visible course experimentation

Evidence still needed. A campus-wide developmental sequence, comparable independent assessments, and long-term learning and relationship evidence

University of Southern California

Strongest public evidence. Broad AI enablement, academic-integrity infrastructure, and protected knowledge boundaries in at least one professional-school policy

Evidence still needed. Consistency across programs, independent verification, privacy implementation, and measured learning

SELECTED INSTITUTIONAL CASES · CONTINUED

Formation and governance must meet

California Institute of Technology

Strongest public evidence. Teaching guidance and an assessment culture capable of protecting live reasoning and demonstrated competence

Evidence still needed. A complete public student-facing governance framework and comparable evidence of AI's educational effects

Loyola Marymount University

Strongest public evidence. The clearest reviewed effort to connect AI literacy with mission, ethics, privacy, access, education, and the humanities

Evidence still needed. Common protected assessments, implementation records, participant evidence, and durable outcomes

Thomas Aquinas College

Strongest public evidence. A unified curriculum of original texts, discussion, proof, laboratory work, writing, and public defense, alongside an AI and human-flourishing initiative

Evidence still needed. A complete public student-facing AI framework covering disclosure, approved uses, privacy, verification, and required unaided work

Why there is no single ranking

Different institutions lead in different dimensions. Some provide stronger formation and protected independent work. Others provide more mature AI infrastructure, privacy controls, student-use evidence, or mission-based ethics. A single numerical ranking would hide these differences and would exceed the available outcome evidence. [6] [7] [8] [9] [10] [11] [13]



Thomas Aquinas College in Santa Paula, one of the California institutions examined for educational purpose and AI governance. Courtesy of Thomas Aquinas College.

CURRENT DETERMINATION

The benefit worth asking for

The choice is not simply whether a school uses AI. It is what the use is for, who remains responsible, which mental acts the learner still performs and what evidence will show an enduring benefit. The most promising path in the reviewed research is bounded assistance connected to accountable teaching and independent demonstration.

California's public systems and selected institutions have published useful principles, teaching guidance, privacy rules, integrity standards, and experimental programs. Some also protect live reasoning, oral defense, original writing, or human-centered formation. The reviewed public record does not yet demonstrate institution-wide durable improvements in independent learning, judgment, mental fitness, or wellness caused by artificial intelligence.

Artificial intelligence should be treated as a bounded educational aid, not as evidence of learning. Schools and universities should protect the mental acts their courses are intended to develop, provide age- and purpose-specific rules, verify understanding after

assistance is removed, and preserve accountable human relationships. A use that improves speed, satisfaction, or the submitted product may be valuable, but it should be described as an efficiency gain unless independent learning also improves.

The strongest institutional model will combine a clear account of human formation with responsible technical governance. Formation without an explicit AI policy leaves teachers, families, and students without practical boundaries. AI infrastructure without an account of educational purpose can normalize use before anyone proves what students retain.

The public-source investigation is sufficiently developed for publication. Implementation and outcomes remain unverified where supporting public evidence is unavailable. No final institutional ranking is justified, and no K through 12 student-facing AI condition is authorized by this project. Future findings should change the article when stronger evidence warrants a change.



Illustration: learners compare field notes with what they observe beside a stream, applying knowledge in a new setting and explaining their conclusions to one another.

A READER'S EVIDENCE GUIDE

Different claims require different proof

A policy, a successful assignment and a lasting capability answer different questions. The links between them have to be demonstrated.

PUBLISHED POLICY

What an institution says it requires

Read the operative document, its authority, scope and date. Publication alone does not show that classrooms follow it.

IMMEDIATE RESULT

What happened during or soon after a task

Inspect the comparison, outcome and conditions. A stronger assisted product may be valuable without showing retained learning.

DURABLE RESULT

What the learner retains or transfers later

Look for unaided performance after time has passed, appropriate comparison and subgroup results. Do not infer it from satisfaction or speed.

The panels separate three kinds of evidence: a published rule, an immediate result and learning that lasts. Each answers a different question and needs its own supporting evidence.

Documented practice

A course, assignment, program, or implementation record shows an activity

Reported experience

An identified participant or institution describes an experience

Context

Family, wellness, COVID, access, disability, or prior condition affects interpretation

Unresolved evidence

The public record does not answer the question and further investigation is recommended

RESEARCH RECORD

What this edition can and cannot establish

This is an editorial review edition dated October 14, 2026. It reorganizes the published education investigation and its linked research. The original content records a September 13 research cutoff and displays September 14 on the main page. The public-evidence-only method was revised September 25. Later dated material remains in its linked thematic record. Selected central study methods were checked for editorial accuracy on October 14; this is not a new comprehensive empirical, institutional or legal review.

GIB evaluates publicly available records, research, audits, regulatory filings and documented outcomes. GIB does not send surveys or questionnaires, conduct interviews, or solicit responses from organizations or individuals. Assessments distinguish published policies and claims from demonstrated results. Where sufficient evidence is unavailable, the finding remains “Insufficient public evidence to assess.” Published surveys and studies conducted by others may be used with their limitations stated. Missing information is not proof of failure or wrongdoing.

Scope of the institutional comparison

The ten institutional cases do not represent every California school or university. Public webpages may omit relevant internal practices, and published language may not predict classroom behavior. The investigation therefore separates mission, policy,

practice, immediate performance, durable learning, context, and missing evidence.

How to interpret the evidence

Research findings report study results. Institutional statements describe published policies or claims. Editorial recommendations propose what should happen. These are different kinds of evidence and analysis.

The [Phase Four Empirical Evidence Review](#) assesses research support by topic as Strong, Moderate, Preliminary or Insufficient. These are GIB assessments, not independent certification or a grade for this article as a whole.

Changes in the public record

The published investigation describes monthly review of the official sites of the ten California institutions. That description is preserved as part of the research method; this editorial edition has not independently verified the current schedule or a fresh monitoring run. A website change can establish a change in the public record. It cannot by itself show classroom implementation or improved learning.

The public change log should record the detection date, institution, education level, official source, prior published position, new published position, significance, and verification status. Routine website edits should not be reported as policy changes.

PUBLIC-EVIDENCE REVIEW CRITERIA

What would make the findings stronger

The earlier investigation organized its remaining gaps by the people or offices best placed to understand them. The categories below are retained as criteria for assessing public evidence. They are not requests for participation. Implementation records, independent learning measures, accessibility and privacy protections, and unsuccessful results would make future conclusions more informative.

Institutions

Operative policies, authority, approved accounts, contracts, implementation records, educator preparation, assessment design, and outcome evidence

Educators

Representative assignments, actual AI rules, workload, enforcement, human-only work, observed benefits, failures, and student explanations

Parents

Notice, consent, alternatives, expectations, ability to influence decisions, family burden, and observed changes

Students

Actual uses, disclosure, confidence, dependency, attention, authorship, relationships, and authorized unaided comparison

Privacy and technology staff

Product configuration, data classification, vendor access, retention, deletion, security, and incidents

Accessibility and wellness staff

Accommodations, equivalent routes, support needs, stress, attention, belonging, and limits on interpretation

Researchers

Comparable groups, preregistered outcomes, baseline condition, delayed measures, transfer, subgroup analysis, and unsuccessful results

TERMS USED IN THIS INVESTIGATION

A reader's glossary

Authorship

Responsibility for the work's claims, reasoning, sources and expression; disclosure identifies material assistance.

Bounded assistance

Help limited by the task, sources, permissions and teaching objective, with the learner's responsibility preserved.

Cognitive offloading

Using an external aid to carry part of a mental task. Whether it helps learning depends on which task is being learned.

Durable result

Unaided retention or transfer measured after time has passed.

Formation

The deliberate development of knowledge, habits, judgment, character and responsible relationships.

Guardrails

Constraints on how a tool responds or is used. Their presence is not, by itself, proof of educational benefit.

Human agency

The person's capacity to understand, choose and take responsibility for action.

Immediate result

Performance measured during or soon after an activity; it may not establish retained or transferable learning.

Independent performance

Demonstrating the intended knowledge or skill after the relevant assistance has been removed.

Mental fitness

A nonclinical educational term used here for attention, recall, judgment, perseverance, self-regulation and responsible participation.

Protected assessment

A planned way to observe the learner's own relevant capability, with fair conditions and appropriate accommodations.

Retention

What a learner can still recall or do after time has passed.

Retrieval practice

Bringing information to mind, instead of only rereading it or obtaining it from a tool.

Source-bounded research

Working from a defined set of materials that the learner can inspect and trace.

Transfer

Applying an idea or skill to a new problem or setting, rather than reproducing a practiced answer.

Working paper

Research circulated before or outside a completed journal-publication process; status and study limitations should be reported.

PRESERVED SOURCE REGISTER

Institutions and educational purpose

The numbered links below retain the original main investigation's source register. The additional MIT research-question link remains in the foundations section. Source 24 is background context, not evidence of an AI effect.

01 California Department of Education Guidance for Safe and Effective AI Use State K through 12 guidance

<https://www.cde.ca.gov/ci/pl/documents/aiguide.pdf>

02 California Department of Education Model AI Policy Model policy and implementation fields

<https://www.cde.ca.gov/ci/pl/aipolicy.asp>

03 California Community Colleges AI System Human capacity governance and system initiatives

<https://ai.cccco.edu/>

04 California State University AI Commons System access resources and training

<https://genai.calstate.edu/>

05 University of California Academic Senate AI Report Agency faculty authority governance and trustworthiness

https://senate.universityofcalifornia.edu/_files/reports/council-chair-to-senate-divisions-senate-ai-workgroup-report.pdf

06 UC Berkeley AI Hub Campus guidance tools and teaching resources

<https://ai.berkeley.edu/>

07 Stanford Understanding AI Literacy Human-centered literacy framework

<https://teachingcommons.stanford.edu/teaching-guides/artificial-intelligence-teaching-guide/understanding-ai-literacy>

08 UCLA Generative AI Guide Research citation and student guidance

<https://guides.library.ucla.edu/generative-ai>

09 USC Artificial Intelligence Resources Institutional AI resources

<https://ai.usc.edu/>

10 Caltech Teaching in the Age of AI Course and assessment guidance

<https://teach.caltech.edu/resources/teaching-in-the-age-of-ai>

11 Loyola Marymount University AI in Teaching and Learning Mission teaching privacy and access

<https://its.lmu.edu/ai/teaching/>

12 Mother of Divine Grace School Curriculum Overview K through 12 classical developmental curriculum

<https://modg.org/curriculum/overview>

PRESERVED SOURCE REGISTER · CONTINUED

Learning, authorship and wellness

13 Thomas Aquinas College AI and Human Flourishing Formation and public AI inquiry

<https://www.thomasaquinas.edu/ai>

14 Harvard-Westlake Curriculum Independent secondary curriculum

<https://www.hw.com/academics/about-the-curriculum>

15 High Tech High Project-based K through 12 model

<https://www.hightechhigh.org/>

16 Bastani and colleagues Generative AI Without Guardrails Can Harm Learning Randomized high-school mathematics study

<https://www.pnas.org/doi/10.1073/pnas.2422633122>

17 Kestin and colleagues AI Tutoring Outperforms In Class Active Learning Peer-reviewed college physics experiment

<https://doi.org/10.1038/s41598-025-97652-6>

18 Wang and colleagues Tutor CoPilot Randomized field trial of AI support for human tutors

<https://edworkingpapers.com/ai24-1054>

19 Risko and Gilbert Cognitive Offloading Review of external tools and mental effort

<https://doi.org/10.1016/j.tics.2016.07.002>

20 Karpicke and Roediger The Critical Importance of Retrieval for Learning Controlled delayed-recall evidence

<https://pubmed.ncbi.nlm.nih.gov/18276894/>

21 Stanford SCALE Evidence Base on AI in K through 12 2026 evidence synthesis and research gaps

<https://scale.stanford.edu/sites/default/files/The%20Evidence%20Base%20on%20AI%20in%20K-12%20Report.pdf>

22 Liang and colleagues GPT Detectors Are Biased Against Non Native English Writers Academic-integrity caution

<https://pubmed.ncbi.nlm.nih.gov/37521038/>

23 Centers for Disease Control and Prevention Youth Mental Health During COVID 19 Pandemic context

https://archive.cdc.gov/www_cdc_gov/media/releases/2022/p0331-youth-mental-health-covid-19.html

24 National Center for Education Statistics 2024 Mathematics Results National achievement context

https://www.nationsreportcard.gov/reports/mathematics/2024/g4_8/national-trends/

THEMATIC READING ROUTES

Read the full investigation

This magazine gives the main argument. The links below open the fuller online investigation, organized by topic. Select a teal heading to read its detailed discussion, evidence or institutional cases. Use Complete project documents for the unabridged research files and Consolidated sources to inspect the underlying references.

Foundations

Classical and modern traditions, developmental stages, knowledge, character and practical judgment.

The modern education system

Institutional hierarchy, curriculum, teachers, parents, alternatives and responsibility.

Influence and funding

How money and authority can reach classroom practice; a relationship alone does not establish improper conduct.

AI and learning

The full learning-evidence review and tool-suitability analysis, including age, privacy and access limits.

The Synthetic You

Present capabilities, prospective imitation, authorship, fair review and the limitations of detection.

Verification standard

Protected performance, retention, transfer, fairness, accessibility and research safeguards.

California findings

Full system comparisons and ten institutional case records.

COVID, family and wellness

Baseline conditions and the limits of claims about learning, technology and wellbeing.

Regulations and oversight

A separately dated tracker. Its policy record is not a new legal review in this magazine.

Public evidence review

Current method, unresolved public evidence and criteria for assessing new records.

Complete project documents

All 22 unabridged web editions and original Word files.

Consolidated sources

The wider cross-investigation register, with original source roles and links.

RESEARCH PRESERVED

The full investigation remains available

The 22 complete project documents remain available in their unabridged web editions and original Word files. They preserve the full research, institutional findings, evidence reviews, verification methods and historical pilot materials. Older descriptions of outreach, participant collection or proposed studies are historical proposals, superseded by the September 25, 2026 public-evidence-only policy. They do not imply that participation occurred or authorize new student-facing activity.

Foundational research master

The Well Formed Mind Research Master

10,492 words · Complete web edition · Original Word file

Review framework

California Education and Artificial Intelligence Institutional Review Guide

5,388 words · Complete web edition · Original Word file

Evidence reviews and scorecards

Phase One Evidence Review

2,000 words · Complete web edition · Original Word file

Phase Two Evidence Review

2,826 words · Complete web edition · Original Word file

Phase Three Course Evidence Review

2,568 words · Complete web edition · Original Word file

Phase Four Empirical Evidence Review

3,113 words · Complete web edition · Original Word file

Phase Five Integrated Institution Scorecard

4,415 words · Complete web edition · Original Word file

Phase Six Verification Investigation Design

2,412 words · Complete web edition · Original Word file

Phase Seven Public Evidence Archive

1,895 words · Complete web edition · Original Word file

Verification and pilot materials

Phase Eight Institution Case Files

4,710 words · Complete web edition · Original Word file

Phase Nine Evidence Ledger

3,224 words · Complete web edition · Original Word file

RESEARCH PRESERVED · CONTINUED

Protocols, cases and supplemental audits

Phase Ten Protected Human Performance Protocol

3,420 words · Complete web edition · Original Word file

Phase Eleven Prospective Pilot Scorecard

3,515 words · Complete web edition · Original Word file

Phase Twelve Pilot Task Packets

3,772 words · Complete web edition · Original Word file

Phase Thirteen Pilot Packet Dry Run

4,401 words · Complete web edition · Original Word file

Phase Fourteen Educator Tabletop Review

3,613 words · Complete web edition · Original Word file

Phase Fifteen Preimplementation Determination

2,781 words · Complete web edition · Original Word file

Phase Sixteen Local Review Execution Binder

2,292 words · Complete web edition · Original Word file

Supplemental institutional audits

California Public Education Systems Artificial Intelligence and Human Formation Audit

2,381 words · Complete web edition · Original Word file

California Campus Level Artificial Intelligence and Human Formation Audit

3,949 words · Complete web edition · Original Word file

California Institutional Artificial Intelligence Readiness and Human Formation Audit

4,169 words · Complete web edition · Original Word file

Artificial Intelligence Tool Suitability Matrix for a Well Formed Mind

4,178 words · Complete web edition · Original Word file

FURTHER READING

A shared concern for human capability

Where Society Meets AI Super-intelligence

Explains model behavior, power and authorship. This investigation tests the learner's independent capability.

AI for Human Flourishing

Uses verified learning benefits in a wider public-value assessment; the education-specific evidence and limits remain here.

California Government Accountability

Examines public authority, spending and oversight, with educational formation considered in the affected institutions.

U.S. Leadership in Artificial Intelligence

Connects human capability with national objectives, implementation and strategic constraints.

A link supplies context and a method for comparison. It does not transfer a conclusion, establish causation or provide independent corroboration.

Formal investigation title: Education and the Well Formed Mind in the Age of Artificial Intelligence.

Project, Research, Analysis, Editorial by Patrick Russell; assisted by ChatGPT (OpenAI).

© 2026 Patrick Russell / Global Intelligence Brief. All rights reserved.