ALCUIN

Project

The Alcuin AI Literacy Framework

Standards for student learning from kindergarten through graduation · September 2026

Thinking is the curriculum.

The Alcuin AI Literacy Framework describes what students should understand about artificial intelligence and how they should use that understanding in school and beyond. It gives schools a sequence of learning expectations, shared classroom norms, and a rubric for examining student use of AI.

All eight strands and 37 standards are drafted. The framework has not yet been field-tested. External verification of the crosswalk is in progress, and the Foundations crosswalk remains to be completed.

The graduate

An AI-literate graduate can use AI effectively and judge its limitations. The graduate understands, in general terms, how a language model produces an answer from patterns in data. They can explain why a convincing answer may be wrong without attributing thought or intention to the system.

In practical work, the graduate chooses a suitable model, supplies the information it needs, and reviews its output. They may use AI to gather information, draft, test, organize, or translate. They check important claims against independent sources and give more attention to errors that could have serious consequences. Claims that cannot be verified are removed or qualified.

This judgment depends on knowledge. A student who knows history can notice a misplaced event in a summary; a student who can calculate can recognize an implausible total. Graduates therefore need a substantial base of knowledge and skills they can use without assistance. They should also understand why some schoolwork must remain their own. Writing a thesis, for example, develops capacities that students do not acquire simply by reviewing one produced for them.

The graduate notices when a chatbot agrees too readily. Before using an unfamiliar app, they ask what it collects, what the provider retains, and how the provider benefits from its use. They understand that having a friend's photograph or a family story does not give them permission to share it with an AI service.

They take responsibility for work submitted under their name. They do not present machine-generated work solely as their own, but rather as a collaboration. They do not use AI to harm others. They do not accept consequential decisions about other people without accountable human judgment.

The graduate can also participate in public decisions about AI. They understand that a feed reflects choices about what to rank and recommend. If a school board is considering an AI product, they can ask what it is intended to do, what data informed its development, and what evidence would show whether it works.

These expectations apply across pathways after high school. A graduate entering a welding program may use AI to draft a bid, then check the figures before submitting it. They should understand what scheduling software collects about workers and seek reliable confirmation when a supplier's chatbot

gives a specification on which a weld will depend. The setting changes; the need for knowledge, verification, and responsible judgment remains.

Purpose and scope

This is a standards framework for schools planning AI literacy instruction. It includes eight strands, 37 coded standards, expectations across four developmental stages, twelve classroom norms, and the Student Use Check. The companion Alcuin Curriculum provides the full standards, including observable indicators, classroom examples, clarifications, and boundaries. The summaries in this volume should be read alongside that full text.

Schools will need lessons, assessments, teacher preparation, and local policies to implement the framework. Those materials are developed separately. Tool selection and academic-integrity policies remain local responsibilities.

Why Schools Need AI Literacy Standards

More than half of American teenagers report having used AI chatbots for help with schoolwork (Pew Research Center, 2026). Schools have often responded by establishing rules about permitted tools and acceptable use. As of December 2024, about a third of public schools had a written policy on student AI use (NCES). Such policies establish necessary boundaries, but they do not by themselves teach students how to judge a system or its output.

Instruction also needs to extend beyond warnings about cheating, privacy, and bias, or demonstrations of efficient prompting. Students need to understand why systems make errors, how to check their work, and how regular use may affect the skills students are developing.

In this framework, AI literacy is the capacity to think, create, judge, and act with AI while understanding the system, the purpose of its use, and the work being delegated. It includes deciding when to work without AI. Students need enough subject knowledge to evaluate what a tool returns, which is why the framework begins with Foundations and uses teacher-mediated, largely device-free instruction in the elementary years.

Two questions guide decisions about use. Will AI help complete the task well? What will its use mean for the student's learning? Other frameworks reviewed for this document address the first question in several ways. AILit asks students to consider the nature of the task; UNESCO asks whether the use serves a justified aim; and the SEE Framework cautions against routine cognitive offloading.

Alcuin gives particular attention to the second question through a developmental sequence of observable expectations. An assignment may be intended to diagnose understanding, demonstrate mastery, or provide practice. In each case, its educational purpose concerns what the student knows

or can do. A completed essay is evidence of learning only if the student has done the work the assignment was designed to develop or assess.

Standards make these expectations teachable. They identify what teachers can observe, how performance develops over time, and what falls outside the framework's scope.

The draft also responds to state and international developments. PISA will assess media and AI literacy in 2029, and this framework includes a crosswalk to AILit, which contributes to that assessment. Arizona advises districts to adopt AI literacy curriculum and professional development tied to state academic standards. Utah's framework directs schools to teach how AI works and when and how to use it; HB 273 establishes a district policy requirement. Utah's framework does not provide student standards, while Arizona refers districts to a district-level set of grade-band progressions. Alcuin offers schools a proposed sequence of student expectations to support this work.

Development and Evidence

Alcuin was developed independently and then compared with published frameworks. In August 2026, the review examined the structure and student competencies in AILit (OECD and European Commission, 2026), UNESCO's AI Competency Framework for Students (2024), Digital Promise's AI literacy framework (2024), and AI for Education's SEE Framework (2026). The review also considered Utah's P-12 framework (March 2026), Arizona's GenAI guidance v26.01 (June 2026), and UNESCO's companion framework for teachers.

The comparison identified fourteen recurring topics, all addressed by Alcuin. It also identified limited treatment of context, memory, model choice, and agentic systems, along with the absence of framework-wide student assessment instruments and subject-level guidance in the seven documents reviewed. These findings informed the standards and companion materials.

Several design choices differ from those sources. Alcuin assigns grade-band expectations, although AILit does not assign its proficiency levels to grades. Alcuin provides observable indicators where UNESCO uses broader competencies. It also treats system evaluation as a general student competency while leaving specialist architecture design to elective computer science. The appendix records these decisions and the supporting page references.

This is a practitioner-developed framework informed by comparative analysis and classroom experience. It is not a research study, and it has not been validated against student outcomes. Comparative findings apply to the documents reviewed. The crosswalk remains subject to external verification.

Classroom Norms

The norms give teachers and students a common language for everyday practice: twelve norms for students working with AI, written for the classroom wall. Each draws on one or more strands; they are not a second set of standards or a one-to-one index to them.

The full text of the twelve norms is part of the Alcuin Classroom poster line and ships with the materials schools purchase.

Organization of the Framework

Each strand addresses a different part of AI literacy. The first three establish the knowledge and judgment needed for the rest of the framework.

STRAND	FOCUS	GUIDING QUESTION
1 Foundations	Knowledge and skills students can use independently	What can I do without AI?
2 Thinking With AI	The effects of use on the learner	How is using AI affecting my learning and judgment?
3 Truth & Evidence	The reliability of output and sources	Is this accurate, and how can I check?
4 The Machine	How systems work	What does this system do, and how does it work?
5 Working With AI	Effective use for a task	Will AI help me do this work well?
6 Data & Ethics	Responsibilities to other people	Whose information is involved, who could be harmed, and who is responsible?
7 AI & the World	Effects on communities and institutions	How is AI shaping what my community sees and does?
8 Shaping AI	Evaluation and improvement of systems	How well does this system work, and who should decide?

Working With AI and Thinking With AI should be considered together. A tool may help produce a good result while taking over work the student needs to practice. Either the demands of the task or its learning purpose may be a reason to leave AI out.

Standards and indicators

Each strand contains four or five standards identified by a code, such as TE.2. Each standard has lettered indicators in the companion curriculum. The strand defines the area of learning, the standard states the outcome, and the indicators identify components teachers can teach and observe.

The opening statement of each standard describes the fully developed competency. The four stages describe progress toward it. Grade bands indicate where that progress is generally expected within a K-12 program.

STAGE	TYPICAL GRADE BAND	DESCRIPTION
Emerging	K-2	Recognizes an idea in shared examples led by the teacher.
Guided	3-5	Applies the idea with support.
Independent	6-8	Applies the knowledge or skill without instructional scaffolding.
Anticipating	9-12	Uses the judgment in planning and carrying out complex work, before a problem requires correction.

A student's starting point depends on demonstrated learning, not age alone. A learner may be at different stages across standards. An adult new to AI, for example, may need introductory instruction about context windows while already evaluating sources independently.

The stages are cumulative unless a standard states otherwise. Indicators collectively define the competency; they are not a checklist that must be demonstrated in full on every assignment. By Anticipating, students should draw on the full competency, although an individual assessment may sample only part of it. Anticipating describes earlier and more independent use of judgment, rather than simply harder work.

A small number of Extension standards begin after prerequisite learning provided elsewhere in the curriculum. Their progressions begin at the first band for which this framework sets an expectation.

Access to AI tools

K-2 and grades 3-5 are readiness bands. Learning is device-free and teacher-mediated, with much of the instruction drawn from Foundations. In the upper elementary grades this includes examining output the teacher has generated, which students compare against a text, an observation, or another source. Grades 6-8 introduce supervised use where product age requirements and school approval allow, with students accounting for what they asked, what the system returned, and what they did with it. Students without eligible access continue from teacher-provided material, developing the same habits. Grades 9-12 include direct use under school policy.

Schools may introduce direct use earlier where their educational judgment and legal requirements allow it. An Independent-stage expectation that requires a tool applies only when the learner has authorized access under both the provider's terms and the school's legal requirements. In current U.S. settings, that is often around age thirteen. The framework assumes access to at least one licensed tool governed by a data processing agreement and configurable by teachers.

Evidence of learning

A student's plan, reflection, or log must be assessed against other evidence, such as an earlier declaration, a prior version, or a live sample of work. A reflective account is never sufficient on its own.

The framework sometimes uses words such as "sees" or "did" to describe a system's behavior. These are shorthand for what it can process or produce. They do not attribute thought, knowledge, or intention to the system.

Core and Extension

A standard is Core if other standards depend on it or if its absence exposes students to immediate harm. Extension standards develop the additional craft, civic understanding, and system judgment described in the graduate profile.

Of the 37 standards, 29 are Core and eight are Extension. Core establishes the minimum for a program claiming alignment and provides a starting point for adoption. Extension identifies work needed to complete the program; it is not optional enrichment. Designations follow the published criteria, so their distribution varies by strand.

A Core-only program prepares students to check AI against their own knowledge, use tools, verify output, protect their learning, and act responsibly toward others. It leaves substantial civic learning unfinished. Three Extension standards, AW.2 through AW.4, address persuasion, employment, and governance. A high school program intended to achieve the graduate profile is expected to teach them and should plan when it will move beyond Core.

One assignment may provide evidence for several standards. For example, a senior review of an actual AI system could address S.1, S.2, S.3, D.5, and AW.4. Each standard would still be assessed against its own expectations. This is an available performance task, not a required capstone.

The Standards in Grid Form

The following summaries retain the framework's standard codes and Core or Extension designations. Each gives the overall expectation followed by the four stages. Full indicators, examples, clarifications, and boundaries remain in the companion Alcuin Curriculum; these summaries do not override them.

Strand 1 · Foundations

Students develop the knowledge and skills needed to work independently and evaluate assistance. No standard in this strand requires an AI tool at any grade band.

F.1 Apply prior knowledge · Core

Draw on knowledge held in memory to judge new claims.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Identifies a false statement about a topic the class has studied.
3-5 Guided	States what they know about a topic before reading and identifies information in the text that contradicts it.
6-8 Independent	Evaluates an unfamiliar claim against prior knowledge and explains whether that knowledge is sufficient to settle the question.
9-12 Anticipating	Before starting a task, distinguishes claims they could check from their own knowledge from those they could not.

F.2 Demonstrate skills without assistance · Core

Complete the core work of coursework without assistance and demonstrate those skills when asked.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Completes a short task involving a core skill independently.
3-5 Guided	Works through a multistep task on paper and marks where they became stuck before asking for help.
6-8 Independent	Produces a complete draft, solution, or argument independently and identifies its weakest part.
9-12 Anticipating	Maintains current evidence of independent performance in each core skill and uses it to plan practice.

F.3 Interpret the task · Core

Restate a task, identify what an adequate answer must include, and recognize when an answer addresses a different question.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Explains the teacher's question in their own words.
3-5 Guided	Sorts questions by the kind of response they require and matches answers to the questions they address.
6-8 Independent	Restates an assignment's requirements before starting and checks the completed work against them.
9-12 Anticipating	Examines a poorly framed question, including its assumptions and criteria for an answer, and revises it.

F.4 Support and reconsider a position · Core

Take a position, support it with reasons, and defend or revise it in response to challenge, keeping a record of the reasoning.

STAGE	STUDENT EXPECTATION
K-2 Emerging	States a claim and gives one reason for it.
3-5 Guided	Responds to one round of disagreement by engaging with the other person's reason.
6-8 Independent	Responds to a prepared counterargument and revises the position when warranted, explaining what changed.
9-12 Anticipating	Sustains a position through extended challenge and records the reasons for the final judgment.

Strand 2 · Thinking With AI

Students consider how AI use affects their learning, judgment, and relationships.

TW.1 Evaluate what is delegated and learned · Core

Identify work delegated to AI, weigh the benefits against what the student gives up, and explain what was learned through the session.

STAGE	STUDENT EXPECTATION
K-2 Emerging	In an unplugged activity, distinguishes help that supports doing a task from help that does it for them.
3-5 Guided	Identifies, with support, which parts of the work the machine completed and which parts the student completed.
6-8 Independent	States at the outset which work will remain their own, compares the closing reflection with that declaration, and demonstrates a sample without the tool.
9-12 Anticipating	Evaluates delegation in unfamiliar work with several parts, makes a plan, and defends it with evidence.

TW.2 Protect work needed for learning · Core

Identify tasks whose educational purpose depends on doing the work, choose to complete them without AI, and explain the decision.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that the class completes some tasks without a computer for a learning purpose.
3-5 Guided	Identifies, with support, whether doing a task is itself the purpose and works accordingly.
6-8 Independent	When the choice is open, decides which work to complete without AI and provides process evidence of that choice.
9-12 Anticipating	Identifies the work that must remain their own in a major project and completes it as planned.

TW.3 Test thinking through challenge · Core

Use AI to challenge reasoning and recognize when a system is agreeing without adequately checking.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Notices when a machine changes its answer after the class changes the question.
3-5 Guided	Asks for a reason someone might disagree and explains whether it changes their view.
6-8 Independent	Independently completes the adversarial review routine and retains the paired records from before and after the review.
9-12 Anticipating	Incorporates adversarial review into their working method without prompting and defends their response in a live discussion.

TW.4 Recognize features that encourage continued use · Core

Identify features designed to extend engagement or simulate a relationship, and explain what the system simulates without possessing.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Identifies a feature in an app that encourages the user to keep playing.
3-5 Guided	Identifies one feature in a classroom tool that encourages continued use and explains its purpose.
6-8 Independent	Audits a tool they use and checks the feature inventory against the app itself.
9-12 Anticipating	Predicts which engagement features an unfamiliar tool may contain, checks those predictions, and explains how companion bots simulate a relationship.

TW.5 Make considered choices about personal authorship · Extension

Recognize the relational meaning of authorship and decide when a message should be written in the student's own voice.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Recognizes the difference between receiving a handmade card and a computer-generated card.
3-5 Guided	Identifies, with support, classroom messages that should come directly from a person and explains why.
6-8 Independent	Writes personal messages themselves, using AI beforehand if needed to clarify what they want to say.
9-12 Anticipating	Makes and defends authorship decisions in contexts observable at school and interprets polished messages from others charitably.

Strand 3 · Truth and Evidence

Students evaluate the accuracy of AI output and the reliability of the sources used to support it.

TE.1 Recognize the risk of confident errors · Core

Anticipate plausible but false output, assess fabrication risk before generating content, and identify claims that need checking.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Finds a false statement in an AI answer read by the teacher and explains how they recognized it.
3-5 Guided	Sorts AI statements into those known to be correct, known to be incorrect, and needing verification.
6-8 Independent	Explains how text prediction can produce a fabricated citation and marks checkable claims before relying on output.
9-12 Anticipating	Identifies parts of a task likely to produce fabricated information and plans verification before generating content.

TE.2 Verify claims before relying on them · Core

Check AI-generated claims against independent sources and justify accepting, revising, or rejecting them.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Suggests one way the class could check a claim in an AI answer.
3-5 Guided	Compares an AI answer with a provided source and explains which account to trust.
6-8 Independent	Checks a claim against two sources, including one located independently, and records the result.
9-12 Anticipating	Includes a verification plan in the project proposal, carries it out, documents results, and resolves discrepancies.

TE.3 Investigate sources through lateral reading · Core

Evaluate unfamiliar claims or sources by consulting independent information about them.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that reading material has a creator and asks who made it.
3-5 Guided	Finds, with support, who is behind a source and reports one relevant fact.
6-8 Independent	Investigates a source through separate searches and compares two independent accounts.
9-12 Anticipating	Builds lateral-reading checks into the research plan, completes them, and documents the inquiry.

TE.4 Assess the origin and authenticity of media · Core

Use provenance and context to assess whether an image, audio recording, or video may be synthetic.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that computers can make pictures of events that never happened.
3-5 Guided	Understands that visual clues alone cannot establish authenticity and asks where an image first appeared.
6-8 Independent	Checks the provenance of a widely circulated image or clip.
9-12 Anticipating	Identifies synthetic-media risks in advance, establishes an authentication procedure, and applies it to a disputed item.

TE.5 Document and disclose AI contributions · Core

Record how AI contributed to the work, distinguish the student's contribution, and disclose use in the required format.

STAGE	STUDENT EXPECTATION
K-2 Emerging	States when the class used a computer's help to make something.
3-5 Guided	Adds a plain-language sentence to completed work explaining what AI helped with.
6-8 Independent	Keeps an AI-use record and writes a disclosure in the class format.
9-12 Anticipating	Plans disclosure at the start of a project and produces a complete account they can explain and defend.

Strand 4 · The Machine

Students learn how AI systems produce output, what information they use, and where their limits lie.

M.1 Explain generative prediction · Core

Explain how a generative system produces output and why the same prompt can produce different answers.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that the computer makes predictions from many examples.
3-5 Guided	Explains that the system predicts words from patterns and observes different answers to the same question.
6-8 Independent	Explains prediction in their own words well enough to account for fabricated output.
9-12 Anticipating	Accounts for variable output, fabrication risk, and sensitivity to prompts when planning a task.

M.2 Explain the role of training data · Core

Describe where training data comes from, whose work contributed to the system, and why those origins matter.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that the computer learned from material people made.
3-5 Guided	Explains that answers draw on patterns learned from people's work and names one kind of training material.
6-8 Independent	Explains how missing training data can lead to gaps in output and tests one prediction about such a gap.
9-12 Anticipating	Identifies reasons a model's training may be incomplete and plans how to address the resulting limitations.

M.3 Identify available context · Core

State what information a system has available during a task and what it does not have.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that the computer can use only the information available to it.
3-5 Guided	Considers whether the system had the necessary information when an answer was wrong.
6-8 Independent	Identifies what a tool can access and checks rather than assuming.
9-12 Anticipating	Plans around context limits, including what must be supplied and what carries over between sessions.

M.4 Recognize AI systems and supervise actions · Core

Recognize AI within familiar tools and exercise judgment about systems that take actions.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that some computer functions operate without a direct request.
3-5 Guided	Identifies AI in classroom tools other than chatbots.
6-8 Independent	Classifies systems in a provided log and infers the objectives their behavior serves.
9-12 Anticipating	Recognizes an agentic system before using it and explains the oversight it requires.

M.5 Consider resource use · Extension

Describe the resources AI systems consume and consider those costs when deciding how to use them.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that AI runs on computers in large buildings that use electricity and water.
3-5 Guided	Traces a query from the classroom to a data center and names two resources used.
6-8 Independent	Uses the fact sheet to explain the difference between training costs and usage costs.
9-12 Anticipating	Uses sourced, dated figures to consider resource costs in an actual tool decision.

Strand 5 · Working With AI

Students learn to select, direct, and supervise AI tools for suitable tasks.

W.1 Frame a clear request · Core

Define the task's goal, constraints, audience, and expected form so a system can assist appropriately.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Helps the class state exactly what it wants the tool to do.
3-5 Guided	Writes a request that includes the assignment's requirements, with support.
6-8 Independent	Includes the goal, constraints, and quality criteria in a request, and diagnoses missing information when a request fails.
9-12 Anticipating	Frames complex requests with several constraints and plans for likely failures.

W.2 Supply relevant material · Core

Provide the sources and other information a task requires, including material maintained in the tool's standing setup.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Observes in a teacher demonstration that a tool can answer better when it has the story being discussed.
3-5 Guided	Includes the necessary material in a request, with support.
6-8 Independent	Supplies suitable material without prompting, checks the output, and explains what information the standing setup provides.
9-12 Anticipating	Plans the information needed for a sustained project and maintains the setup that makes it available.

W.3 Evaluate and revise output · Core

Treat initial output as a draft, evaluate it, direct revisions, and decide when it meets the task's requirements.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Helps the class decide whether an answer is adequate or needs another attempt.
3-5 Guided	Identifies one specific change an output needs, with support.
6-8 Independent	Evaluates output against the assignment's criteria and either directs revision or gives a reason for accepting it.
9-12 Anticipating	Uses criteria established in advance to guide revision and explains decisions to retain material or stop revising.

W.4 Decide whether and where AI is useful · Core

Decide whether a task calls for AI, choose appropriate tools, and deliberately allocate the work.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Helps the teacher and class decide whether a computer would help with a task.
3-5 Guided	Explains, with support, whether AI suits a task and which part it could assist.
6-8 Independent	Makes a suitability decision where use is permitted and identifies the work assigned to the tool.
9-12 Anticipating	Plans the allocation of a task with several parts and explains the decisions after completing it.

W.5 Supervise sustained work · Extension

Set checkpoints, identify departures from the plan, and remain accountable for AI contributions throughout a project.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Observes the teacher pause a demonstration project to check the machine's work.
3-5 Guided	Checks at a teacher-set checkpoint whether the AI-assisted work still matches the plan.
6-8 Independent	Sets checkpoints and answers an unannounced question about work in progress without opening the tool.
9-12 Anticipating	Plans supervision for a system that can act independently and remains accountable for the completed work.

Strand 6 · Data and Ethics

Students protect information, examine bias, and understand responsibility for AI use and decisions.

D.1 Examine data practices · Core

Find out what an AI tool collects, stores, uses for training, and shares.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that an app may retain information shared with it and that the class must be careful.
3-5 Guided	Identifies, with support, information an app collects and locates one privacy setting.
6-8 Independent	Uses public documentation to answer the four data-practice questions for a tool they use.
9-12 Anticipating	Evaluates an unfamiliar tool's data practices before use and explains the tradeoffs.

D.2 Protect other people’s information and work · Core

Respect other people's control over their private information, likeness, and unpublished work.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that the class does not share other people's private information with a computer.
3-5 Guided	Asks permission, with support, before using someone else's material with a tool.
6-8 Independent	Applies the rule without prompting in realistic situations and explains why possession does not establish permission.
9-12 Anticipating	Anticipates consent questions in a project and includes permission requests in the plan.

D.3 Examine sources of bias · Core

Explain how bias can enter a system through data, design, and use.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that unfair examples can contribute to unfair computer responses.
3-5 Guided	Identifies, with support, a pattern in a tool's answers.
6-8 Independent	Identifies plausible sources of bias, tests what can be tested, and distinguishes findings from what remains unknown.
9-12 Anticipating	Identifies bias risks in a proposed use and designs checks to address them.

D.4 Prevent and respond to harmful use · Core

Apply prohibitions on harmful use, respond appropriately when targeted, and reason through less clear cases.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that computers must not be used to hurt others and tells an adult when this happens.
3-5 Guided	Explains the rules for the grade band and recognizes clear violations.
6-8 Independent	Applies the full prohibitions and the proportionality question to realistic cases, including ambiguous ones.
9-12 Anticipating	Reasons about proportionality in unfamiliar cases and explains the applicable limits.

D.5 Question automated judgments about people · Core

Explain why people remain accountable for decisions and ask informed questions about systems used to evaluate students or others.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that people remain responsible for decisions about people, even when a computer helps.
3-5 Guided	Explains, with support, that people must check computer-generated judgments about people.
6-8 Independent	Explains how detection or proctoring tools can be wrong and who may be affected.
9-12 Anticipating	Prepares for an actual automated evaluation by identifying the system, relevant questions, and a route for raising concerns.

Strand 7 · AI and the World

Students examine how AI influences information, employment, and public decisions.

AW.1 Explain how recommendations shape information · Core

Explain the objectives of recommendation systems and consider how different feeds affect a shared community.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that a computer selects which video plays next.
3-5 Guided	Explains, with support, that apps may select content to keep people watching and that people chose this objective.
6-8 Independent	Explains ranking and its objective, tracing how two users' behavior can produce different feeds.
9-12 Anticipating	Evaluates design alternatives and community effects and writes an audit of their own information consumption.

AW.2 Analyze persuasion at scale · Extension

Explain how generated content and precise targeting affect persuasion, and use that understanding to interpret the information environment.

STAGE	STUDENT EXPECTATION
K-2 Emerging	No separate expectation in this framework.
3-5 Guided	No separate expectation in this framework.
6-8 Independent	Explains how targeting and content generation work together, identifies likely examples, and recognizes ordinary identifiable advertising.
9-12 Anticipating	Predicts how a campaign might target someone like them and analyzes an actual campaign operating at scale.

The elementary years rely on prerequisite media and advertising literacy taught elsewhere in the curriculum. This is a design decision, not a claim that every school already provides that instruction.

AW.3 Examine changes in work and employment · Extension

Describe how AI affects work and use evidence about those changes in personal planning.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that machines have long helped people work and that people remain responsible for decisions.
3-5 Guided	Names jobs in which computers perform part of the work and identifies what the person still does.
6-8 Independent	Distinguishes a task from a job using examples and identifies parts of a familiar occupation affected by AI.
9-12 Anticipating	Uses sources to examine changes in a field they may enter and identifies skills they plan to develop in response.

AW.4 Participate in decisions about AI governance · Extension

Describe how AI is governed, from district policy to national regulation, and practice relevant forms of civic participation.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Understands that school computer rules are made by people.
3-5 Guided	Identifies, with support, who makes technology rules at familiar levels.
6-8 Independent	Explains the district's AI policy and identifies one question it leaves unresolved.
9-12 Anticipating	Uses sources to follow a current governance question and prepares a civic contribution to the standard required for submission.

Strand 8 · Shaping AI

Students investigate systems, test their performance, and propose evidence-based changes.

S.1 Review system documentation · Core

Investigate an AI system's intended purpose, users, and stated limitations.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Through teacher-led examples, understands that people and organizations make tools for particular reasons.
3-5 Guided	Finds, with support, a tool's stated purpose and compares it with how the tool is used.
6-8 Independent	Identifies the purpose, users, and limitations described in documentation, along with information the documentation does not provide.
9-12 Anticipating	Independently consults an unfamiliar system's documentation and evaluates it against the proposed use.

S.2 Evaluate performance against criteria · Core

Evaluate a system against stated criteria using test cases and users' experience.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Helps the class check whether a tool performs as promised in a teacher-run demonstration.
3-5 Guided	Tests a tool against one class-selected question and reports the results accurately, with support.
6-8 Independent	Establishes criteria before designing a test and reports results another student could reproduce.
9-12 Anticipating	Raises evaluation questions without prompting, conducts the evaluation, and makes a recommendation supported by evidence.

S.3 Decide whether a system is appropriate · Extension

Determine whether a problem calls for AI and specify conditions under which the proposed system would fail or should be stopped.

STAGE	STUDENT EXPECTATION
K-2 Emerging	Recognizes that a class problem might be addressed by a computer, a person, or a rule, and helps choose an approach.
3-5 Guided	Generates and compares three approaches, identifying one drawback and one affected person for each.
6-8 Independent	Defines a real school problem and identifies what an AI approach would require.
9-12 Anticipating	Produces a complete scope document that includes failure conditions and reasons to stop the system.

S.4 Specify an improvement · Extension

Translate a complaint or evaluation finding into a specification someone responsible for the tool could use.

STAGE	STUDENT EXPECTATION
K-2 Emerging	With the teacher's help, states the situation and the change they would like.
3-5 Guided	Replaces a general complaint with a specific description of what should happen in a given situation.
6-8 Independent	Writes an evidence-based specification for a tool they have tested or used.
9-12 Anticipating	Produces a specification that considers the developer's tradeoffs and is written for someone responsible for the tool.

AI Literacy and Critical Thinking

Students need regular opportunities to think through difficult work themselves. Habitually delegating that work may weaken capacities they no longer practice. The research is still developing, but the concern is sufficient to justify protecting independent work within the Thinking With AI strand.

AI can also provide material for explicit critical-thinking instruction. Teachers can use generated arguments, explanations, and source analyses for students to examine. A flawed example gives the class something concrete to question, while a structured exchange can provide repeated practice responding to counterarguments. Teachers remain responsible for selecting the material and guiding the instruction.

This use is reflected in the standards. In TE.1, young children identify an error in a teacher-presented AI example. In TW.3, students practice evaluating counterarguments and explaining whether the challenge changes their position. Generated work can also let students practice direct criticism of a text before applying that precision, with appropriate care, to classmates' work.

What the research supports

A 2026 systematic review of 67 higher-education studies found the strongest critical- and creative-thinking outcomes when ChatGPT was used in scaffolded, inquiry-oriented instruction. Passive or weakly structured use was more often associated with cognitive offloading and lower engagement (Li, Cui, and Hagedorn, 2026).

Xi, Zhang, and Wang (2026) reported stronger academic performance and reflective thinking among university students using a Socratic AI agent than among those using a non-Socratic agent. In a separate study, Namaziandost and Teng (2026) reported greater gains in critical thinking, epistemic

beliefs, and argumentation under a Socratic-AI condition than under generic-chatbot and no-AI conditions.

Most evidence still comes from postsecondary settings. K-12 evidence is more limited. It includes Tang and colleagues' (2026) study of critical questioning in one tenth-grade English classroom and Kao, Grant, and Woltering's (2025) randomized trial in tenth-grade science, which had not passed peer review at the draft's review date.

These studies do not establish that AI-supported instruction is generally better than strong instruction without AI. The framework's position is that deliberately structured use can make some kinds of critical-thinking practice easier to provide repeatedly across a classroom. Field validation will examine whether implementation reliably provides those conditions.

Definition and curriculum connections

The framework uses critical thinking in the Paul-Elder tradition. Paul and Elder describe eight elements of reasoning and intellectual standards for evaluating them. Repeated practice is intended to develop habits such as intellectual humility, autonomy, and fairmindedness (Paul and Elder, 2019).

Critical thinking is addressed across the strands rather than in a separate strand. The following mappings distinguish direct instruction from partial support and learning assigned to other subject areas.

ELEMENT OF THOUGHT	TREATMENT IN THE FRAMEWORK
Purpose	W.1 addresses the task's purpose, S.1 the system's purpose, and TW.1 the learning purpose.
Question at issue	W.1 and S.3 require students to examine the task or problem being addressed.
Information	TE.2, TE.3, M.2, M.3, and W.2 address evidence, sources, training data, and available context.
Interpretation and inference	D.3 and M.4 require inferences from observed behavior and tests of those inferences. TE.4 requires evidence-based judgments about authenticity. General analysis of inference remains in subject-area argument instruction.
Concepts	M.1 teaches the concept of generative prediction. Identifying and examining the concepts that organize an argument remains in subject-area instruction.
Assumptions	M.2 and M.3 examine assumptions about training and available information. S.3 examines assumptions in problem framing. General analysis of unstated premises remains in subject-area argument instruction.
Implications and consequences	S.3, TE.2, AW.1, and TW.1 address the consequences of systems, claims, information selection, and delegation.
Point of view	TW.3, TE.3, and AW.1 address counterarguments, independent accounts, and differences in information exposure.

INTELLECTUAL STANDARD	TREATMENT IN THE FRAMEWORK
Clarity	W.1 and S.4 require clear requests and specifications.
Accuracy	TE.1, TE.2, and TE.3 address errors and verification.
Precision	S.4, W.1, and W.3 require specific requests, criteria, and improvements.
Relevance	Addressed indirectly through material selection in W.2, task suitability in W.4, and important claims in TE.2. It is not taught as a general intellectual standard.
Depth	Supported by difficult tradeoffs in S.3 and strong counterarguments in TW.3, but not defined or assessed as a general standard.
Breadth	S.3, TW.3, TE.3, and AW.1 require consideration of alternatives and perspectives.
Logic	TW.3 includes testing through counterargument. General argument structure remains in subject-area instruction.
Significance	S.4 and TE.2 require attention to consequential findings and claims.
Fairness	TW.3 requires a strong opposing case; D.5 examines automated judgments; D.3 examines bias and affected groups.

INTELLECTUAL TRAIT	TREATMENT IN THE FRAMEWORK
Intellectual humility	D.3 distinguishes findings from inference; M.2 avoids claims about undisclosed training data; TE.2 requires removing or qualifying unverified claims; TE.4 recognizes limits of detection.
Intellectual courage	D.5 supports questioning automated judgments, while D.2 and D.4 address acting when others are at risk. Broader engagement with contested ideas belongs elsewhere in the curriculum.
Intellectual empathy	Partially addressed through student, teacher, and administrator perspectives in D.5, different feeds in AW.1, and charitable interpretation of messages in TW.5.
Intellectual autonomy	TW.1, TW.2, and W.5 require students to retain and exercise judgment.
Intellectual integrity	TE.5 and TW.5 address disclosure and authorship. W.5 requires accountability for completed work.
Intellectual perseverance	TW.2 protects productive struggle. Broader perseverance through difficult work is a responsibility of the surrounding curriculum.
Confidence in reason	Implicit in testing claims, giving reasons, and revising judgments in response to evidence. It is not taught as a separate disposition.
Fairmindedness	TW.3 gives equal credit to a reasoned change of mind and a reasoned decision to retain a position. D.3 asks which groups a system's bias affects.

Alcuin is not a complete critical-thinking curriculum. General argument analysis, including concepts, unstated premises, and the relationship between reasons and conclusions, remains a subject-area responsibility. The mappings should not be read as claims of complete coverage.

The developmental stages reflect the gradual development of these habits. Teachers initially supply the questions and guide the response. Students then use them with support, apply them independently, and eventually bring them into new work without prompting. Schools can use the mappings when connecting Alcuin to existing graduate profiles and critical-thinking requirements.

The Student Use Check

The Student Use Check assesses how AI was used in producing work. The quality of the work itself is assessed separately. The five dimensions are the decision to use AI, direction of the tool, verification, disclosure, and the work the student completed themselves.

Score the decision only when students had a real choice. If the assignment required AI, mark that dimension not applicable and assess the remaining dimensions. If AI was prohibited, use the school's academic-integrity policy rather than this rubric.

Open-choice settings, such as independent research or writer's workshop, allow students to practice deciding. When a student makes a reasoned choice to work without AI, score only the decision dimension. Evidence of that choice must include observed work or a sample completed without the tool; a finished product cannot establish that AI was absent.

Evidence may include the plan, process records, version history, completed work, observed performance, or a conversation. At least one source must come from outside the student's written account. Students do not need a polished verbal explanation, and a rehearsed explanation is not sufficient evidence of understanding.

Using the rubric

Use the rubric for conferences and feedback until the school has established examples from student work and teachers have practiced scoring them consistently. Shared examples and evidence of scoring reliability are prerequisites for grading.

Report the AI-use judgment separately from the content grade. An essay may meet the writing criteria while failing the disclosure expectations, and both results should remain visible.

Use sampled work for a full review. One complete review per student per term is a practical starting point; scoring every assignment for a large teaching load is not the intended use.

The Developing descriptors include responses given when a teacher asks. That conversation is part of the assessment. Asking a student to explain a passage or work through a step can provide useful evidence of understanding.

Scoring descriptors

The decision

W.4 and TW.1. Apply only when the student had a choice.

LEVEL	DESCRIPTOR
4 Exemplary	Deliberately chose whether to use AI and explains both its value for the task and its effect on learning. The process is consistent with those reasons.
3 Proficient	Made a deliberate choice and explains it in terms of the task.
2 Developing	Gives a reason when asked, but the plan and process contain no evidence that the choice was considered.
1 Beginning	Used or declined AI without a reason they can explain.

The direction

W.1 through W.3.

LEVEL	DESCRIPTOR
4 Exemplary	Provided relevant material and constraints, evaluated output against clear criteria, and directed revision where needed. The process demonstrates purposeful control.
3 Proficient	Framed the task clearly, evaluated output, and revised instructions or material when needed.
2 Developing	Framed the task incompletely or accepted output with little evidence of evaluation.
1 Beginning	Delegated the task with little direction and used the returned output.

The verification

TE.1 and TE.2.

LEVEL	DESCRIPTOR
4 Exemplary	Matched checking to the consequences of error and documented it during the work. Removed or flagged claims that could not be verified.
3 Proficient	Checked claims essential to the work against a source independent of the tool and kept a record.
2 Developing	Checked isolated details or relied on asking the same tool again.
1 Beginning	Used the output without examining it.

The disclosure

TE.5.

LEVEL	DESCRIPTOR
4 Exemplary	Gives a specific, complete account of what was generated, revised, and independently produced. An informed reader would find no undisclosed contributions.
3 Proficient	Names the tool and its contribution in the required format.
2 Developing	Gives a vague disclosure, such as "I used AI," or leaves gaps.
1 Beginning	Provides no disclosure or an account contradicted by the work.

The student's work

TW.2.

LEVEL	DESCRIPTOR
4 Exemplary	Declared in advance which work would be completed without AI and why. The completed work supports that declaration.
3 Proficient	Provides visible evidence that the student completed the work the assignment was intended to develop.
2 Developing	AI has replaced thinking intended for the student, who cannot explain the relevant passages when asked.
1 Beginning	The product displays capacities the student cannot demonstrate or explain in a sample completed without the tool.

Implementation and Review

Field validation

The framework is published as a working draft. Pilots will examine whether students verify claims, disclose use, and make different delegation decisions after instructional support is removed. Counts of standards covered will not be treated as evidence of learning. Findings will inform public revisions.

Coverage and boundaries

The framework covers generative, recommendation, and agentic systems. Additional system types can be incorporated where the existing strands apply. Architecture design and model training remain outside its scope and belong in elective computer science. Students are expected to investigate and evaluate systems and specify improvements, rather than build scalable architectures.

Elementary instruction prepares students through checking claims, comparing sources, and distinguishing real events from invented representations. It is largely unplugged and does not require young children to use AI independently.

Annual review

The enduring understandings and standard statements are intended to remain useful as tools change. Examples, tool details, and much of Working With AI require annual review against current technology and research. The evidence underlying Thinking With AI should receive particular attention because it is still developing. Each review cycle will be recorded in the version history.

Teacher preparation and assessment

Implementation requires preparation in practices such as lateral verification, model selection, and reading model cards. A teacher prerequisite map, crosswalked to UNESCO's teacher framework, and a professional-development sequence are companion materials.

Schools must establish scoring examples, practice, and reliability before using the Student Use Check summatively. Other assessment instruments are separate companion products.

Core and Extension designations may be reconsidered against the published criteria. Core identifies prerequisites and immediate protections; Extension completes the graduate profile.

Glossary

Agentic AI. Systems that take actions, such as searching, clicking, filing, purchasing, or scheduling, with varying degrees of human oversight.

Artificial intelligence or AI. Systems that perform tasks through computational models, most often learned from patterns in data. This framework includes classifiers, recommendation systems, optimizers, and generative models. Their functions differ, so no single account of how they work applies to all of them.

Bias in AI. Systematic skew in data, behavior, output, or decisions. It can arise through training data, design choices, or use. People can increase or reduce it intentionally or unintentionally.

Context window. The limited amount of material a model can consider during a task. Material outside that window does not inform the answer.

Detection software. Tools that attempt to identify AI-written work. They can produce errors, and research has found uneven false-positive rates across groups of students. A detector result is not sufficient evidence on its own.

Fabrication or hallucination. Plausible, confident output that is false. It is a known failure of generative prediction and does not involve an intention to deceive.

Fine-tuning. Additional training that adjusts an existing model for a particular purpose or customer.

Generative AI. Systems that produce text, images, audio, video, code, or other content through statistical prediction based on patterns learned from training data.

Human in the loop. An arrangement in which a person retains judgment and accountability within an AI-assisted decision.

Large language model or LLM. A generative system trained on large amounts of text that produces language by predicting the next token. LLMs power most chatbots.

Memory in AI tools. A feature that allows a model to use information from earlier sessions. What the model can use later is distinct from what the provider retains. Providers may retain conversations under their policies whether or not a memory feature makes them available in a later session.

Model. The trained system. Differences among models can affect their suitability and performance.

Model card. A provider's description of a system's purpose, development, and limitations. It should be evaluated critically like any other source.

Prompt. The question, instructions, and material the user supplies to the system. Their quality can substantially affect the output.

Provenance. The origin and history of a claim, image, or piece of work, including its earliest appearance and subsequent circulation.

Recommendation system. AI that ranks and selects information, such as feeds, videos, prices, or search results, according to objectives set by its operators. Those objectives may differ from the user's.

Retrieval-augmented generation or RAG. A design that allows a model to consult selected sources before answering. It can improve grounding and accuracy but does not eliminate errors.

Synthetic media including deepfakes. Machine-generated images, audio, or video, including convincing depictions of people or events that did not occur.

System prompt. Standing instructions supplied by a tool's provider or the school. They shape responses and are usually not visible in the conversation.

Token. A small unit of text processed or produced by a language model. Models operate on tokens and learned numerical representations.

Training data. Material from which a model learned patterns. Its sources, the work involved in producing it, and consent for its use may be documented or disputed.

Walled garden. A closed system whose provider controls the content, tools, and data within it. Schools may use such systems to control access and data flows.

Appendix A · Framework Crosswalk

This appendix records the draft relationships between Alcuin and the frameworks reviewed in August 2026. It covers the seven strands drafted at that time. Foundations was added on September 8 and its crosswalk is pending. All mappings await external verification before publication.

Each entry lists the corresponding provisions in six comparison documents. UNESCO's teacher framework informed the wider review but is not a column in the student crosswalk. "None identified" means the review found no substantial counterpart; "partial" indicates limited overlap. Page references are to the comparison documents, not this volume.

For compactness, the entries use AILit for OECD/EC 2026, UNESCO for the 2024 student framework, Digital Promise for its 2024 framework, SEE for its 2026 framework, Arizona for v26.01, and Utah for its 2026 framework.

Truth and Evidence

TE.1. AILit: Knowledge 3.2, fabrication risk (p. 22), and Engage 3, evaluation (p. 29). UNESCO: none identified. Digital Promise: Practice 6, information and mis/disinformation (p. 8). SEE: Stay Critical (p. 15) and evaluating and refining output (p. 23). Arizona: Goal 4 (p. 23). Utah: Principle 3, partial (p. 6), and hallucination glossary entry (p. 15).

TE.2. AILit: Engage 3, accepting, revising, or rejecting output (p. 29). UNESCO: none identified. Digital Promise: Practice 6, reliable evidence (p. 8). SEE: evaluating and refining output (p. 23). Arizona: Goal 4, comparison with evidence and trusted sources (p. 23). Utah: review before use in the overreliance provision (p. 11).

TE.3. AILit: Engage 3, partial (p. 29). UNESCO: none identified. Digital Promise: Practice 6 explicitly names lateral reading (p. 8). SEE: none identified. Arizona: Goal 4, partial. Utah: none identified.

TE.4. AILit: Knowledge 3.2, misrepresentation and manipulation (p. 22). UNESCO: none identified. Digital Promise: Practice 4, evaluation of synthetic content during creation, partial (p. 7). SEE: Act Responsibly, deepfake creation prohibition, partial (p. 16). Arizona: Consideration 2 (p. 12) and Goal 4. Utah: none identified.

TE.5. AILit: Create 4, authenticity and intellectual property, partial (p. 35). UNESCO: none identified. Digital Promise: Practice 4, ethical sourcing and citation (p. 7). SEE: disclosure of GenAI use (p. 19), the closest match identified. Arizona: Goal 3, attribution (p. 23). Utah: Principle 5, credit to sources and tools (p. 7).

The Machine

M.1. AILit: Knowledge 1.1-1.3 (p. 21). UNESCO: AI foundations (p. 24). Digital Promise: Practice 1, partial (p. 7). SEE: How GenAI Works (p. 11). Arizona: Goal 1 (p. 23). Utah: Principle 3, how AI works, partial (p. 6).

M.2. AILit: Knowledge 2.1-2.2 and 2.4 (p. 22). UNESCO: AI foundations, partial (p. 24). Digital Promise: Practices 2-3, dataset quality and training data (p. 7). SEE: human choices in development (p. 14). Arizona: Goal 1, data and human design (p. 23). Utah: none identified.

M.3. AILit, UNESCO, and Digital Promise: none identified. SEE: knowledge of context limits, partial (p. 11). Arizona: none identified. Utah: glossary terminology only (pp. 15-16). The review identifies context as an area of limited coverage.

M.4. AILit: Knowledge 1.4 (p. 21) and Engage 1 (p. 29). UNESCO: AI foundations, partial (p. 24). Digital Promise: none identified. SEE: no counterpart within its GenAI-only scope. Arizona: Goal 1, partial. Utah: agentic AI and human oversight (p. 12), plus the glossary (p. 15).

M.5. AILit: Knowledge 1.5 (p. 21) and Engage 5 (p. 29). UNESCO: sustainability and embodied ethics (p. 22). Digital Promise: Practice 5, environmental costs, partial (p. 8). SEE: Act Responsibly (p. 16). Arizona: Consideration 7 (p. 18). Utah: none identified.

Working With AI

W.1. AILit: Computational Thinking skill (p. 24). UNESCO: application skills (p. 24). Digital Promise: Practice 1, prompting (p. 7). SEE: prompting and context-setting (p. 22). Arizona: Goal 2 (p. 23). Utah: Principle 3, how to use AI (p. 6).

W.2. AILit: Collaboration skill, goals and constraints (p. 23). UNESCO: application skills, partial (p. 24). Digital Promise: none identified. SEE: prompting and context-setting, partial (p. 22). Arizona: Goal 2, partial. Utah: none identified. The review identifies sustained context provision as thinly covered.

W.3. AILit: Create 3 (p. 35) and Collaboration, iteration (p. 23). UNESCO: application skills, partial. Digital Promise: Practice 4, output evaluation, partial (p. 7). SEE: evaluating and refining output (p. 23). Arizona: Goal 2, supporting rather than replacing thinking (p. 23). Utah: overreliance warning (p. 11).

W.4. AILit: Manage 1-3 (p. 39). UNESCO: proportionality, partial (p. 22). Digital Promise: Practice 5, suitability (p. 8), and Practice 1, decomposition (p. 7). SEE: deciding when to use GenAI (p. 21). Arizona: Goal 2, when and how (p. 23). Utah: Principle 3, when to use AI (p. 6).

W.5. AILit: Manage 3-4, decomposition and monitoring (p. 39). UNESCO: human determination, partial (p. 23). Digital Promise: none identified. SEE: maintaining human agency, partial (p. 17). Arizona: Goal 5, partial (p. 23). Utah: human oversight (p. 12).

Thinking With AI

TW.1. AILit: Problem Solving skill, partial (p. 24). UNESCO: human-centred mindset, partial (p. 21). Digital Promise: none identified. SEE: avoiding cognitive offload (p. 22), the closest approach identified. Arizona: Goal 2, protecting student thinking (p. 23). Utah: overreliance warning (p. 11) and agentic offloading provision (p. 12).

TW.2. AILit: Responsible attitude, including choosing not to use AI (p. 26). UNESCO and Digital Promise: none identified. SEE: avoiding cognitive offload, partial (p. 22). Arizona: Goal 5, Human Agency and the Human Advantage, partial (p. 23). Utah: none identified.

TW.3. AILit: Critical Thinking and Collaboration skills, including perspectives (p. 23). UNESCO: human-centred mindset, partial (p. 21). Digital Promise: none identified. SEE: Stay Critical (p. 15). Arizona: Goal 4, partial (p. 23). Utah: none identified.

TW.4. AILit: Critical Thinking skill, chatbot wellbeing (p. 23), and Engage 1 Advanced example, engagement design (p. 30). UNESCO and Digital Promise: none identified. SEE: healthy balance (p. 18), Be Intentional (p. 15), and commercial engagement design (p. 14). Arizona: Consideration 6, mental wellbeing (p. 17). Utah: none identified.

TW.5. No substantial counterpart identified in the six student comparison documents.

Data and Ethics

D.1. AILit: Knowledge 2.3, data gathered during use (p. 22), and Collaboration privacy awareness (p. 23). UNESCO: safe and responsible use, consent over data (p. 23). Digital Promise: Practice 3 (p. 7). SEE: data privacy and security (p. 18). Arizona: Consideration 4 (p. 14). Utah: security, privacy, and safety (p. 13).

D.2. AILit: none identified. UNESCO: safe and responsible use, anonymized examples, partial (p. 18). Digital Promise: Practice 3, partial. SEE: protecting privacy, partial. Arizona: Consideration 4, partial. Utah: prohibition on personally identifiable information, including others' data (p. 13). Everyday permission to use another person's material receives limited treatment beyond these provisions.

D.3. AILit: Knowledge 2.5 (p. 22) and Engage 6 (p. 29). UNESCO: non-discrimination (p. 22). Digital Promise: Practice 5 (p. 8). SEE: biased training data (p. 11). Arizona: Consideration 1 (p. 11) and Goal 3 (p. 23). Utah: Principle 1, evaluation for bias (p. 6), and bias glossary entry (p. 15).

D.4. AILit: Self and Social Awareness, including cyberbullying and deepfakes (p. 25). UNESCO: do no harm (p. 22). Digital Promise: Practice 4, responsible sharing (p. 7). SEE: avoiding harmful use and obtaining consent (p. 20), and Act Responsibly (p. 16). Arizona: Considerations 2 and 6 (pp. 12, 17). Utah: prohibited uses (p. 11).

D.5. AILit: Knowledge 4.2, audit and regulation, partial (p. 22). UNESCO: human accountability (p. 21) and entitlement to transparency (p. 23). Digital Promise and SEE: none identified. Arizona: Goal 5, agency and accountability (p. 23). Utah: Principle 6 (p. 7) and detector false-positive warning (p. 11).

AI and the World

AW.1. AILit: Engage 4, recommendations and perspective (p. 29), and Knowledge 4.1 (p. 22). UNESCO: citizenship, partial (p. 21). Digital Promise: none identified. SEE: no counterpart within its GenAI-only scope. Arizona and Utah: none identified.

AW.2. AILit: Knowledge 3.2, manipulation (p. 22), and Self and Social Awareness (p. 25). UNESCO: citizenship and societal effects, partial (p. 21). Digital Promise: Practice 6, bias in reporting (p. 8). SEE: societal impacts (p. 20). Arizona: Consideration 2 (p. 12). Utah: none identified.

AW.3. AILit: Knowledge 4.1, partial (p. 22). UNESCO: citizenship and lifelong learning, partial (p. 21). Digital Promise: none identified. SEE: introductory hiring and access evidence (p. 5). Arizona: State 48 discussion, partial (p. 23). Utah: none identified.

AW.4. AILit: Knowledge 4.2 (p. 22). UNESCO: ethics by design, reviewing regulation (p. 23), and AI divides among countries (p. 22). Digital Promise: Practice 5, access and equity (p. 8). SEE: Act Responsibly, environmental impacts and unequal access, partial (p. 16). Arizona: Consideration 5 (p. 16) and the guidance as an example of governance. Utah: Principles 2 and 7 (pp. 6-7) and the framework as an example of governance.

Shaping AI

S.1. AILit: Shape 1, intended operation, audience, and limitations (p. 43). UNESCO: transparency and explainability, partial (p. 23). Digital Promise: Practice 3, investigating models and methods, partial (p. 7). SEE, Arizona, and Utah: none identified.

S.2. AILit: Shape 2, criteria, test cases, and user feedback (p. 43). UNESCO: iteration and feedback, partial (p. 25). Digital Promise, SEE, and Arizona: none identified. Utah: Principle 7, audit and evaluation, partial (p. 7).

S.3. AILit: Engage 7, ethical principles and values, partial (p. 29). UNESCO: problem scoping, whether AI should be used, and iteration and feedback, including shutdown reasoning (p. 25). Digital Promise: Practice 5, whether and how a tool is suitable (p. 8). SEE, Arizona, and Utah: none identified.

S.4. AILit: Shape 4, improvement for societal benefit (p. 43). UNESCO: iteration and feedback, partial (p. 25). Digital Promise, SEE, Arizona, and Utah: none identified.

Deliberate exclusions

Alcuin does not cover UNESCO's architecture-design expectations in Aspect 4, including scalable architectures and prototypes (pp. 25-26), or AILit's Shape 3 system-design competency (p. 43). These specialist expectations belong in elective computer science. Shaping AI covers problem definition, evaluation, and improvement for all students.

AILit's Basic, Intermediate, and Advanced levels are not assigned to Alcuin's grade bands. AILit explicitly does not tie them to ages or grades (p. 27). Any correspondence is recorded at the standard level. Alcuin's four stages are its own developmental structure.

Decisions about use and learning

The review distinguished deciding whether AI suits a task from deciding whether doing the task is necessary for the student's development. It did not identify an observable standard and developmental progression for the latter question in the six student comparison documents.

FRAMEWORK	TREATMENT OF THE USE DECISION	TREATMENT OF THE LEARNING PURPOSE
AILit	Manage AI 1 considers the nature of the task (p. 39); Responsible attitude includes choosing not to use AI (p. 26).	No explicit rationale, observable standard, or progression identified for preserving the capacities developed by doing the work.
UNESCO Students	Proportionality asks whether a use serves a justified aim (pp. 22-23).	No explicit rationale, observable standard, or progression identified for this purpose.
Digital Promise	Asks whether and how AI is the right tool for the job (p. 8).	No explicit rationale, observable standard, or progression identified for this purpose.
SEE	Addresses when to use GenAI and avoiding cognitive offload (pp. 15-22).	Addresses the risk most directly, framed as effective use. No observable standard or developmental progression identified for the educational purpose of doing the work.
Arizona	Goal 2 calls for supporting rather than replacing thinking (p. 23).	Suggests the rationale but provides no observable standard or progression for it.
Utah	Principle 3 includes when to use AI (p. 6).	No explicit rationale, observable standard, or progression identified for this purpose.

The additional August 2026 review of all five AI4K12 progression charts found one item addressing the human side of the use decision: a grades 9-12 debate prompt about language models and homework. The 2025 CSTA/AI4K12 learning priorities include The Choice to Use AI, framed around decision quality. Digital Promise's March 2026 scan of 33 frameworks also identifies a broader lack of clarity about how competencies develop from childhood through adolescence.

Appendix B · References and Evidence Notes

The references and review statuses below are retained from the September 2026 working draft. They document the draft's evidence base and its remaining publication work. They do not represent a new source-verification exercise for this edition.

Framework references

OECD and European Union. 2026. *Empowering Learners for the Age of AI: An AI Literacy Framework for Primary and Secondary Education.* OECD Publishing. doi:10.1787/65cd27d4-en. Approved by the PISA Governing Board on April 8, 2026; contributes to the PISA 2029 Media and AI Literacy assessment.

UNESCO. 2024. *AI Competency Framework for Students.* UNESCO Digital Library, ark 48223/pf0000391105.

UNESCO. 2024. *AI Competency Framework for Teachers.* UNESCO Digital Library, ark 48223/pf0000391104.

Digital Promise. 2024. *AI Literacy: A Framework to Understand, Evaluate, and Use Emerging Technology.* ERIC ED671235.

AI for Education. 2026. *The SEE Framework: A Practical Guide to Building Generative AI Literacy.* aiforeducation.io.

Arizona Institute for Education and the Economy, Northern Arizona University. 2026. *Generative AI in K-12 Education: Guidance for Arizona Schools, v26.01.* azk12.ai.

Utah State Board of Education. 2026. *Artificial Intelligence Framework for Utah P-12 Education.* schools.utah.gov.

These seven documents were consulted in August 2026. The source draft records links as checked at consultation. Complete live links and consultation details are to be assembled from the source register for publication.

Student use and school implementation

Teenage use of chatbots. Pew Research Center, *How Teens Use and View AI*, published February 24, 2026. The survey was fielded September-October 2025 among 1,458 teenagers aged 13-17. It reports that 54% had used AI chatbots for schoolwork help and 10% did all or most schoolwork with chatbot help. The draft marks this verified. The claim concerns having used chatbots, not all forms of AI or frequency of use.

Employer attitudes toward AI skills. LinkedIn's *AI and the Global Economy* (April 2025) restates the Microsoft/LinkedIn 2024 Work Trend Index, a survey of 31,000 workers across 31 markets. It reports that 66% of leaders said they would not hire someone without AI literacy skills. The draft marks the corrected statement verified. It describes business decision-makers' stated attitudes in 2024, not observed hiring behavior.

Written school policies. NCES School Pulse Panel data from December 2024 report written student-AI policies in 31% of public schools, including 43% of high schools and 27% of elementary schools. The draft marks these figures verified. An earlier claim that just over half of schools had policies by 2025 was removed because the cited EdWeek evidence measured educators, not schools, and did not support the claim.

Access to teacher training. RAND and Diliberti and colleagues (2025) report that 67% of low-poverty districts provided AI teacher training, compared with 39% of high-poverty districts. The draft rates this evidence as solid.

PISA assessment. The OECD/EU AILit framework was approved and declassified by the PISA Governing Board on April 8, 2026. The draft records verification against the local PDF. The framework contributes to the 2029 Media and AI Literacy assessment; that relationship should not be overstated. doi:10.1787/65cd27d4-en.

Utah instruction and policy. Principle 3 of Utah's P-12 framework addresses how AI works and when and how to use it. The district policy requirement comes from HB 273, Classroom Technology Amendments (2026), effective July 1, 2026. The source draft marks both claims verified and distinguishes the law from the framework. It also records that the framework contains no student competency standards.

Arizona guidance. Arizona's GenAI guidance v26.01 (June 2026), p. 22, advises adoption of AI literacy curriculum tied to state standards. Page 23 refers to Chandler USD for grade-band progressions. The source draft records verification against the guidance.

Utah detector warning. Utah's framework warns on p. 11 that detectors can produce false positives, including for non-native speakers. The source draft records verification against the local PDF.

Learning and engagement

Cognitive offloading. The draft draws on Lodge and Loble (2026, UTS expert review); Burns and colleagues (2026, Brookings policy analysis); Bastani and colleagues (randomized trial); Kosmyna and colleagues (preprint with formal criticism); and Gerlich (2025, corrected correlational study). Evidence that repeated delegation erodes unpracticed capacities remains early and contested. It is treated as a reasoned basis for the Thinking With AI strand, not settled science. Morrin and colleagues (2026) was removed because it had been miscited and concerns AI-related delusions.

Chatbot agreeableness and dependence. Sources include Sharma and colleagues (ICLR 2024), Cheng and colleagues (Science, 2026), Kim, Xie, and Yang (arXiv:2512.15117, 2025), and Malmqvist (2025). The draft marks as verified the qualified claim that chatbots can become overly agreeable in ways that promote dependence and repeated use. Sharma addresses preference-training pressures, which should not be presented as proof of corporate intent. Kim and colleagues address relational style and emotional closeness among adolescents.

Features that extend engagement. Sources include Eyal (2014); Monge Roffarello and colleagues (CHI 2023); the Wall Street Journal Facebook Files (2021); California v. Meta (N.D. Cal., 2023, involving more than 41 attorneys general); and the MDL 3047 order (2023). Streaks, continuation prompts, and autoplay are described as designed to extend use. Claims that they were designed to cause addiction remain contested. This evidence is distinct from research on chatbot agreeableness.

Definition of critical thinking. Paul, Richard, and Linda Elder.

1. *The Miniature Guide to Critical Thinking Concepts and Tools*, eighth edition. Rowman & Littlefield. The source draft records verification on August 27, 2026, against criticalthinking.org pages 497, 527, and 528. The Guide's expanded intellectual-standards list includes significance; the website's

eight-item list does not. The expanded list in this framework is attributed to the Guide. **Instructional position.** The framework proposes that carefully structured AI use can make repeated critical-thinking practice easier to provide. It does not claim general superiority over instruction without AI. The draft's comparative claim was removed, and field validation is intended to test whether implementation creates the proposed instructional conditions. **Structured and Socratic use.** Li, Cui, and Hagedorn (2026), *Computers and Education: Artificial Intelligence*, 10, 100571, review 67 higher-education studies. Their review includes harms associated with unstructured use; associations should not be restated as established treatment effects. Xi, Zhang, and Wang (2026), *Computers & Education*, 241, 105494, report a randomized study in Chinese higher education comparing Socratic and non-Socratic agents. Namaziandost and Teng (2026), *International Journal of Applied Linguistics*, doi:10.1111/ijal.70268, report a quasi-experiment among Iranian male EFL learners. These are separate studies with distinct populations and results. Tang and colleagues (2026), *Thinking Skills and Creativity*, 59, 102043, report action research in one tenth-grade classroom; it supports instructional design claims. Kao, Grant, and Woltering (2025), Research Square, doi:10.21203/rs.3.rs-8118546/v1, report a randomized trial in tenth-grade science with a null finding for metacognitive self-regulation. At the draft's review date, it was a preprint without peer review. The source draft records publisher checks on August 27, 2026, and calls for rechecking Kao's publication status during subsequent review. Xi and colleagues are cited as 2026, matching the volume publication date despite a 2025 DOI stem. **Argument analysis.** Assigning general analysis of concepts, assumptions, and logic to English language arts and other subject curricula is a design decision recorded on August 27, 2026. The internal Paul-Elder matrix assessed this coverage as adjacent across the 33 standards then present; the framework now contains 37. The matrix and decision log remain the supporting records for future coverage decisions.

Fairness and protection

AI-text detection. Liang and colleagues, *Patterns* 4(7), 2023, doi:10.1016/j.patter.2023.100779, tested seven detectors available in

1. Average false-positive rates were 61.3% for TOEFL essays and 5.1% for native-English essays. The draft marks the finding verified within that scope. The article is an Opinion piece, concerns short essays, and did not test Turnitin, which disputes English-language-learner bias in its tool. General statements about current detectors require qualification. **Proctoring and disability.** Sources include U.S. Senate inquiry letters (Blumenthal and colleagues, December 2020), EPIC's complaint to the D.C. Attorney General (December 2020), the NDLSA bar-exam report (July 2020), CDT reports (2020 and 2022), the ACM US Technology Policy Committee statement (2023), and Kwapisz and colleagues (arXiv:2511.10826, 2025, preprint, n=493). Concerns about disability-related false flags are documented in inquiries, complaints, and civil-society research. The source draft does not claim an established peer-reviewed finding. **Bias in a résumé screener.** Reuters' 2018 investigation by Jeffrey Dastin documented an experimental Amazon recruiting tool that penalized references to women and downgraded graduates of women's colleges. The tool was

scrapped. The draft marks this verified, with the qualifications that the tool was experimental and Amazon disputes production use. **Removal of nonconsensual intimate imagery.** The draft cites the TAKE IT DOWN Act, P.L. 119-12, signed May 19, 2025, with a platform removal regime in force May 19, 2026, and FTC enforcement. It also cites NCMEC's Take It Down service, introduced in 2023 and covering AI-generated imagery of minors. The source draft marks the statement current as of late May 2026. This supports D.4's reference to removal routes.

Systems and social effects

Environmental resources. The source draft records a completed environmental fact sheet with nine sourced and dated entries covering water, energy, mineral supply chains, training, and inference. Sources include Google's 2025 self-report, Epoch AI (2025), Mistral's 2025 life-cycle assessment, IEA (2025/26), LBNL (2024), Luccioni (FAccT 2024), and Li and colleagues on water use, with a note on disclosure gaps. The fact sheet supports M.5 and requires annual review. Its internal record is framework/factsheet_environment.md.

Interpretability. The claim that systems cannot always be given a complete, human-readable explanation draws on 2024-2026 interpretability literature. It is date-sensitive and requires annual verification.

Training on conversations. Consumer chatbot data practices depend on dated provider policies and account settings. School licenses are governed by their data processing agreements. D.1 distinguishes these contexts; a claim about consumer free tiers should not be generalized to school deployments.

Retention and memory. Dated provider retention policies support the distinction between a memory feature and company retention of conversations. Turning off memory does not by itself establish that a provider retains no conversation data.

Prompt and context quality. Prompting literature and provider documentation support the statement that the quality of supplied instructions and material can substantially affect output.

Employment effects. Task automation examples, including paralegal document review, draw on labor economics research. Claims about whole-job replacement remain contested. AW.3 distinguishes existing changes to tasks from forecasts about occupations and requires dated sources.

Commerce ranking. Platform-specific reporting and litigation records support the statement that ordering can reflect factors beyond relevance. AW.1 uses that qualified description.

Recommendation objectives. Sources include Zhao and colleagues (RecSys 2019, YouTube), Covington and colleagues (RecSys 2016), X's open-source ranker (2023), Instagram engineering (2023), Gomez-Uribe and Hunt (2016, Netflix), and the FTC v. Amazon complaint (2023). The draft marks as verified that ranking objectives can combine engagement, retention, relevance, and revenue. These sources also support the qualified commerce-ranking statement.

Elementary media-literacy prerequisites. State English language arts and media-literacy standards registries inform AW.2's decision to rely on prerequisite instruction in K-5. This is a curriculum design choice, not evidence that provision is universal.

Additional crosswalk evidence

AI4K12 progressions. The source draft records direct reading of all five working-draft charts in August 2026, with detailed reads of Big Ideas 3 and 5 on August 27. The charts date from 2020-2022. The identified use-decision item is 5-B-i, grades 9-12, Unpacked. The review found no developmental treatment of the learning purpose of retaining work. Supporting material is archived in analysis/94.

CSTA and AI4K12 priorities. *AI Learning Priorities for All K-12 Students* (May 2025) was read directly. Its grades 3-5, 6-8, and 9-12 expectations frame The Choice to Use AI through evaluation, debate, and analysis of decision quality. The source draft records verification and an archive in analysis/94.

Digital Promise scan. Digital Promise's March 26, 2026 scan of 33 AI literacy frameworks identifies insufficient explanation of how competencies differ for younger children and adolescents. The source draft records a check against the published post on August 27, 2026.

AILit numbering. The final June 2026 AILit framework places Manage AI 1 in the list on p. 39 and detail on p. 40. The source draft records verification on August 27, 2026, and confirms 19 competencies distributed 7/4/4/4. These references supersede earlier references to Manage 2 for this decision.

Remaining publication work

Complete the Foundations crosswalk and external verification of the mappings. Assemble full citations, live links, and consultation dates from the source register. Retain the evidence qualifications above, and update date-sensitive claims during each review cycle.