ARIP Education System: Enhancing Student Reasoning in the Age of AI

Executive Summary

Generative AI is rapidly entering classrooms, but education systems face a fundamental problem: **the final assignment no longer reliably shows how much thinking the student actually did**. A polished essay, analysis or recommendation may reflect genuine student reasoning, heavy AI assistance, or almost complete cognitive outsourcing. Recent academic studies show that existing approaches largely focus on banning AI, detecting AI-generated text, or providing AI tutors. These remedies, by themselves, do not provide a student with a formal method to optimize reasoning, or produce a governed record for institutions to assess how a student employed AI in the thinking process.

The AI-Augmented Research & Insight Platform ("ARIP") provides students with a reasoning process that is a structured, auditable workflow rather than leaving it hidden inside an AI chat or behind a finished submission. Within ARIP, a student works under an authenticated User ID and progresses an idea or assigned problem through a phased governance lifecycle:

1. **DEFINE** — state clearly what is being tested or argued.
2. **MEASURE** — build the claims, evidence and counter-evidence.
3. **ANALYZE** — evaluate support, contradiction, uncertainty and weaknesses.
4. **IMPROVE** — revise the reasoning where the evidence requires it.
5. **CONTROL/CLOSURE** — record and justify the outcome, along with any remaining uncertainty.

The value of the system is not the sequence alone. **ARIP records the operations performed inside that sequence.** It distinguishes material originated by the student from material introduced by AI; records whether AI suggestions were accepted, rejected or revised; links claims to evidence; preserves the order in which decisions were made; and maintains an append-only history of how the student's analytical position changed. This creates something conventional assignments and AI-detection tools cannot provide: **reasoning provenance**.

An educator is no longer limited to asking, "Did this student use AI?" Instead, the educator can examine:

- what the student originated;
- what AI contributed;
- how the student evaluated AI suggestions;
- whether contradictory evidence was seriously considered;
- whether the student revised a position when warranted;
- what evidence supported the final conclusion; and
- whether the student retained intellectual control over the work.

A conventional chatbot can generate answers, but it does not enforce institutional rules around attribution, sequence, evidence, acceptance, revision and closure. A word processor records a document, not the reasoning that produced it. This is why a dedicated platform is required. AI-detection tools attempt to infer authorship after the fact. ARIP instead **governs the reasoning process as it occurs**. For schools and universities, this creates a potential new model for responsible AI adoption: allow students to use powerful AI tools, but require their reasoning, judgment and participation to remain visible and accountable.

At institutional scale, the same ARIP architecture allows educators to assess reasoning development across assignments and over time, while preserving privacy, teacher judgment and academic-integrity controls. Rather than treating AI as something education must either prohibit or surrender to, ARIP offers a third path:

Use AI — but make the human reasoning observable, governable and assessable.

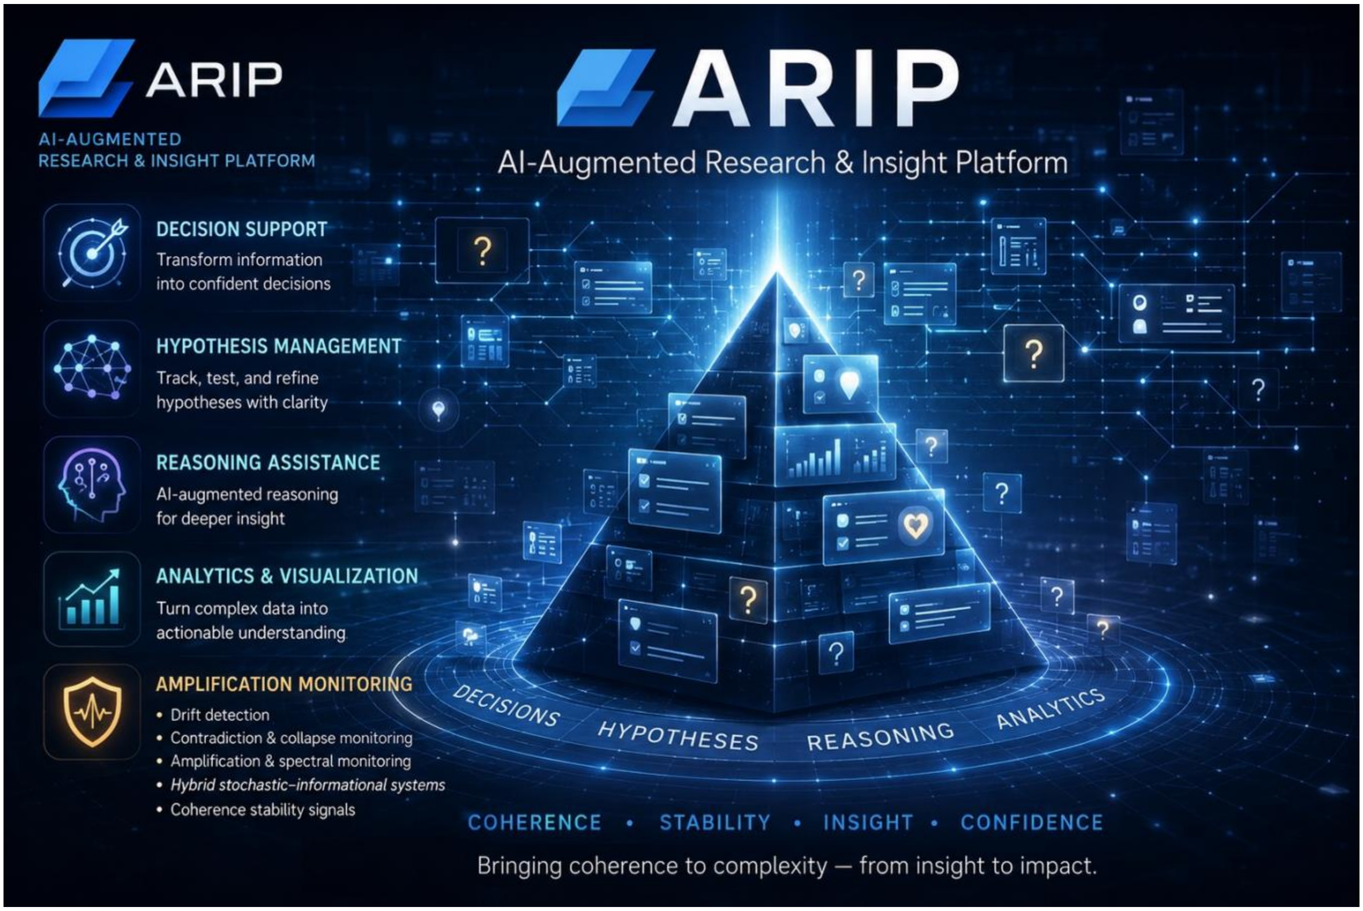

Figure 1. The AI-Augmented Research & Insight Platform (ARIP)

1. The Educational Problem

Generative AI presents an increasingly important challenge for education: students can use powerful AI systems to complete academic work while bypassing some of the cognitive effort through which learning, reasoning and judgment are developed.

The strongest experimental evidence to date concerns not whether students use AI, but how that use is structured. In a randomized field experiment with nearly 1,000 high school mathematics students, Bastani et al. (2025) deployed two GPT-4 tutors: one resembling a standard chatbot interface, and one prompted with learning safeguards. Access to the unrestricted tutor improved performance substantially while the tool was available, but when it was withdrawn those students performed approximately 17 percent worse than peers who had never had access. The safeguarded version, which withheld direct answers and required students to work through the reasoning, largely eliminated that deficit. The finding matters for ARIP because it locates the harm in the mode of interaction rather than in AI itself, and shows that structural constraints on how a student engages with AI can prevent it.

Emerging Canadian research provides evidence of this risk. In a 2025 experimental study conducted in a first-year management course at the University of Toronto Mississauga, students completing AI-supported case analysis retrieved information more quickly but demonstrated reduced peer discussion and lower average performance than students completing traditional group case-study activities. Classroom observations suggested that AI use shifted student behaviour toward information consumption rather than collaborative knowledge construction.

A 2026 York University study of 107 fourth-year civil-engineering students similarly found widespread academic use of generative AI. The researchers concluded that current usage patterns were allowing students to offload cognitive work at points where they would ordinarily spend time engaging with new knowledge and developing deeper understanding.

These Canadian findings are consistent with a wider international literature. Lehmann, Cornelius and Sting (2025), combining observational data from university programming courses with two pre-registered laboratory experiments, found no uniform effect of large language model use on learning. What mattered was usage behaviour: students who substituted AI for their own work by generating solutions covered more topics but understood each one less well, while students who used AI to ask for explanations deepened their understanding. Taken together with Bastani et al., the evidence points consistently in one direction. The educational risk lies in unstructured substitution, and it is addressable by designing how the interaction occurs.

1.1 Why Generative AI Intensifies the Challenge

If foundational knowledge is weak, generative AI can conceal the gap by producing fluent explanations, essays, calculations, and arguments. Traditional assignments often evaluate only the submitted artifact. In that environment, the instructor may be unable to distinguish strong student reasoning supported by AI from weak reasoning replaced by AI. The emerging issue is therefore not simply whether students use AI. It is **how AI is incorporated into the learning process**. Unstructured use can allow students to outsource reasoning and cognitive effort.

The educational opportunity is not simply to detect AI use. It is to redesign the learning environment so that AI assistance remains available while the student's judgment, evidence handling and reasoning remain visible. Appropriately governed use can instead require students to remain intellectually engaged, evaluate AI outputs, confront contradictory evidence, and demonstrate how their own thinking progressed. This creates the central educational opportunity addressed by ARIP: rather than trying to remove AI from student work, design the human–AI interaction so that **reasoning remains visible, active and assessable**.

> ***The shift:*** *The policy question changes from "Did the student use AI?" to "How did the student reason with AI, and what evidence demonstrates meaningful human participation?"*

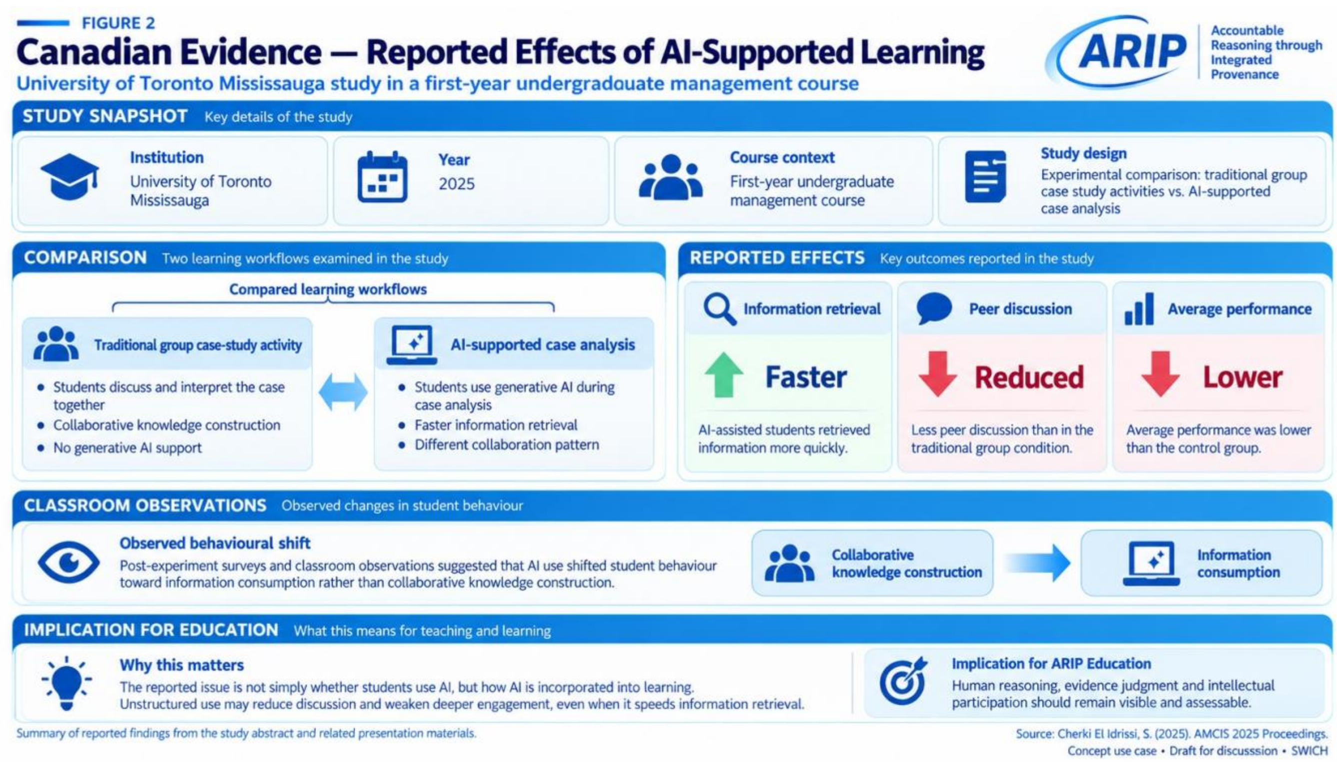

Figure 2. Canadian evidence

1.2 Institutional Recognition of the Problem

Institutions are converging on the same conclusion. In August 2026, MIT's Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training reported that generative AI is making it materially harder to assess student progress and is eroding the implicit contract between instructors and students. Two of its findings bear directly on the ARIP Education model.

First, the committee recommended against relying on AI-detection software, warning that detection invites an arms race with "humanizer" tools and that even low false-positive rates fall disproportionately on non-native English speakers and neurodivergent students. In its place, the committee pointed to process evidence: instructors can require students to work on platforms that capture a history of how work developed and submit that history alongside the work itself. This is an endorsement of the category ARIP occupies rather than the detection category, from an institution with no commercial interest in the outcome.

Second, and more pointedly, the committee identified an evidentiary gap in academic integrity processes. MIT's framework was built on a binary model in which submitted work was either the student's own or someone else's, and generative AI collapses that binary. The committee noted that faculty are uncertain where the line now falls, that the Institute must clarify what evidence would support bringing an integrity case where AI is involved, and that detector output alone is not sufficient for that purpose.

This is precisely the gap a governed reasoning record addresses. ARIP does not attempt to infer whether text was machine-generated. It records, at each material decision point, what was contributed by whom and what the student did with it. That converts an unanswerable question about a finished document into an attributable record of conduct, the kind of evidence integrity processes are designed to weigh.

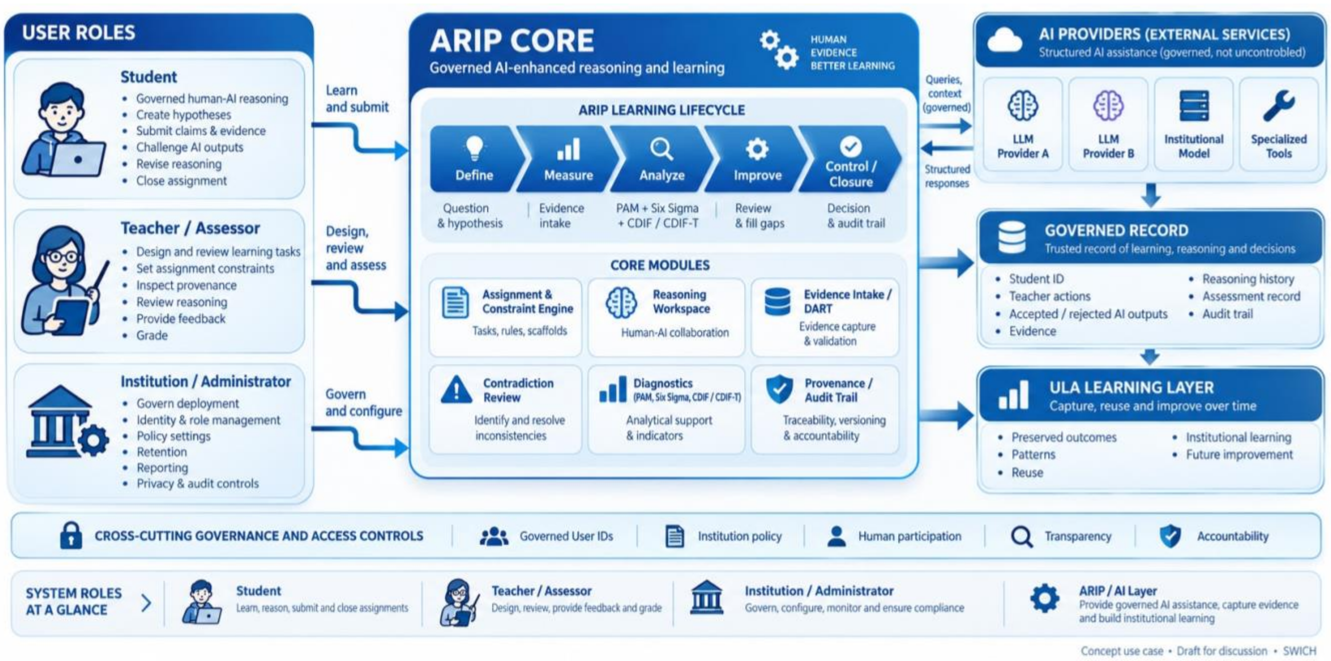

Figure 3. Institution deployment architecture

2. The ARIP Education Module

Within ARIP, a student works through an assignment as a governed sequence rather than a single exchange with a chatbot. At **IDEA/ASSIGNMENT**, the student receives a problem or proposes one, and their initial thinking is captured before any AI expansion. **DEFINE** converts that starting position into an explicit, testable proposition, with its scope, assumptions and intended outcome stated and fixed before evidence is admitted. In **MEASURE**, the student builds the claims and assembles evidence on both sides, recording where each source came from and why it is credible. **ANALYZE** examines what the evidence actually supports, where it is contradicted, and where the reasoning is weakest, with AI used to surface the strongest counter-explanations rather than to supply conclusions.

IMPROVE is where the student responds: revising claims, adding evidence, or changing the proposition where the contradiction warrants it, or deliberately holding the position where it survives challenge. **CONTROL/CLOSURE** closes the assignment with a conclusion, a statement of what changed and why, and an account of the uncertainty that remains.

The stages alone are just a sequence; the value is in what the system records as the student moves through them. Each material contribution is marked with its origin, so work introduced by AI is distinguishable from work originated by the student. AI suggestions do not enter the record automatically — the student must accept, reject or modify each one, and that decision is stored with the actor and the time it was made. Evidence is linked to identifiable sources, revisions preserve what came before rather than overwriting it, and the order of decisions is retained so that an early position cannot be quietly rewritten once the student knows how the analysis ends. The result is a reasoning record an instructor can review directly: what the student framed, what AI offered, what the student did with it, how the position moved under challenge, and whether the final conclusion is proportionate to what the record actually supports.

A school, university, faculty or individual course could deploy a dedicated ARIP environment. Each learner would receive a governed User ID, while instructors would receive assignment-authoring, review and assessment permissions. The institution would define how AI may be used, what activities must be completed, what evidence is admissible, and which elements contribute to assessment.

Role	Primary Function	Illustrative Capabilities
Student	Conduct governed human–AI reasoning	Create hypotheses; submit claims/evidence; challenge AI outputs; revise reasoning; close assignment
Teacher / Assessor	Design and review learning tasks	Set assignment constraints; inspect provenance; review reasoning; provide feedback; grade
Institution / Administrator	Govern deployment	Identity/role management; policy settings; retention; reporting; privacy and audit controls
ARIP / AI Layer	Support structured reasoning	Prompt structured inquiry; surface contradiction; preserve accepted outputs; compute governed diagnostics where appropriate

2.1 The Student Journey

An assignment can be structured around the existing ARIP lifecycle rather than around a single prompt-and-answer exchange. The educational implementation can be adapted for age, subject and course requirements while preserving ARIP's minimum governance floor: origin attribution, explicit student disposition of AI contributions, evidence traceability, authenticated actions, stage integrity, and append-only revision history.

1. IDEA / ASSIGNMENT — Student receives a problem or proposes an idea. Initial thinking is captured before AI expansion.

2. DEFINE — The proposition is converted into a clear, testable hypothesis or question. The student must explain scope, assumptions and intended outcome.

3. MEASURE — Claims and evidence are constructed. The student identifies what supports or weakens the proposition and records source quality.

4. ANALYZE — The student and AI examine support, contradiction, uncertainty and weak points. The system preserves what was proposed and what the student accepted.

5. IMPROVE — The learner responds to counter-evidence, revises claims, adds evidence or changes the hypothesis where warranted.

6. CONTROL / CLOSURE — The student records a final conclusion, explains what changed, and states the reasoning supporting closure.

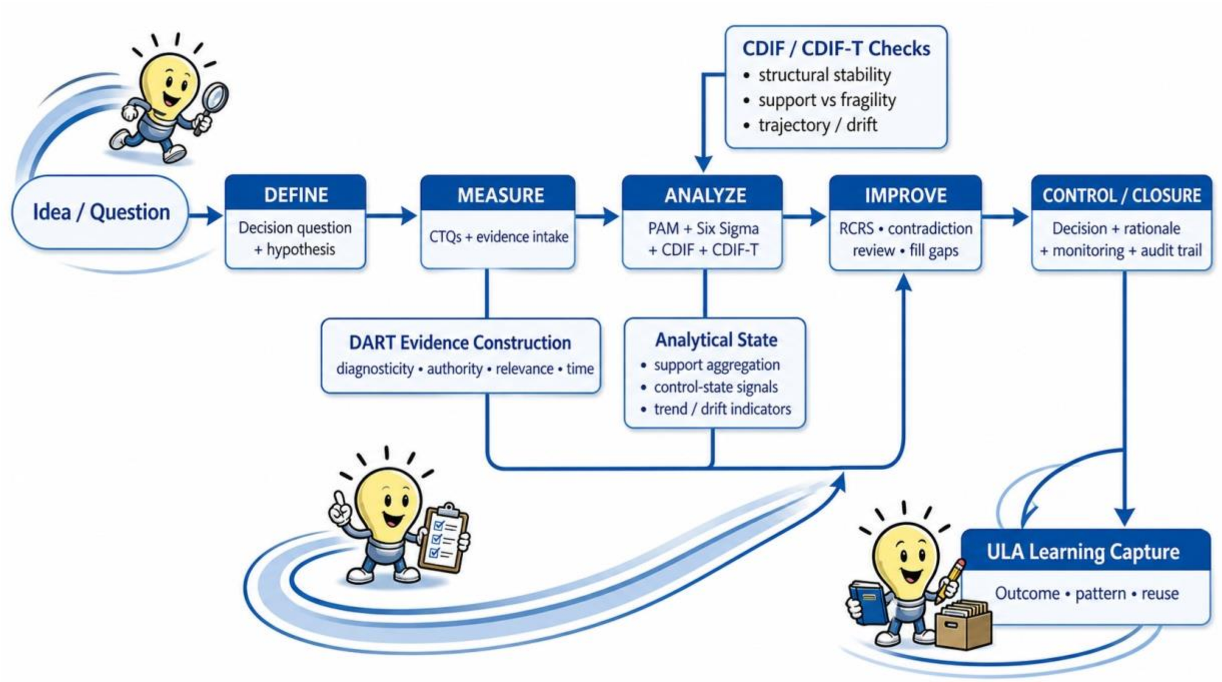

Figure 4. Core ARIP Lifecycle Process

2.2 What the Teacher Can Review

The instructor would not be limited to the final essay, calculation or recommendation. A governed record could expose the reasoning pathway behind the output.

- Student-originated propositions and revisions.
- AI-generated suggestions presented to the student.
- Student acceptance, rejection or modification of AI suggestions.
- Evidence selected, excluded or challenged.
- Contradictions identified and the student's response to them.
- Material changes in the hypothesis or conclusion.
- Final rationale for closure.
- Time-stamped participation and an auditable sequence of interaction.

Figure 5. Teacher review dashboard example

2.3 Assessment Model: Human Contribution and Reasoning Quality

The ARIP Education module supports a teacher-defined rubric that separates subject-matter correctness from reasoning quality and from the quality of human–AI interaction. The goal would not be to calculate a simplistic 'percentage human.' Instead, provenance and interaction data would help the teacher judge whether the learner exercised meaningful intellectual control. A detailed rubric can be found at **Appendix A**.

Assessment Dimension	What the Teacher Looks For	Illustrative Evidence
Idea formation	Student frames or meaningfully shapes the problem	Initial proposition, scoping choices, original questions
Evidence judgment	Student distinguishes stronger from weaker support	Source selection, exclusion rationale, evidence challenges
Contradiction handling	Student engages with information that weakens the preferred answer	Rejected assumptions, revised claims, counter-explanations
AI challenge	Student evaluates rather than automatically accepts AI output	Accept/reject/edit trail, corrective responses
Revision quality	Reasoning changes when evidence warrants it	Version history, hypothesis revision, updated conclusion
Closure quality	Final conclusion is proportionate to the evidence	Final rationale, uncertainty statement, remaining questions

An important design principle: ARIP does not equate more clicks, more text or more disagreement with better reasoning. Teacher judgment remains central. Any automated indicators of student contribution or reasoning quality requires separate educational validation.

2.4 Longitudinal Learning Through ULA

The Unified Learning Architecture (ULA) in ARIP creates an additional educational opportunity. With appropriate privacy, governance and validation, repeated assignments could produce a longitudinal record of how a learner's reasoning develops over time. ULA would observe preserved lifecycle records rather than silently rewriting them.

Potential longitudinal questions could include:

- Does the learner become better at formulating testable propositions?
- Does evidence selection improve in quality and relevance?
- Does the learner become more willing to revise a preferred explanation?
- Does reliance on unchallenged AI suggestions decrease over time?
- Does contradiction handling become more sophisticated?
- Do reasoning patterns transfer across subjects?

These are research questions, not validated metrics in the current ARIP specification. An education deployment therefore begins with transparent qualitative and descriptive indicators, then validates any proposed quantitative scoring against educator judgment and real learning outcomes.

2.5 Minimum Governance Floor

ARIP Education can be adapted for different ages, subjects, course levels, and assessment models; since analytical complexity can vary substantially. A Grade 10 history assignment should not be forced to use the same quantitative machinery as a university investment-analysis course. However, the properties that make ARIP a governed reasoning environment rather than a simple six-step worksheet should not be simplified away.

At minimum, every ARIP Education deployment should preserve the following governance requirements:

- **Origin attribution.** Material contributions should identify whether they originated with the student, AI, instructor/system, or another governed source.

- **Explicit student disposition of AI contributions.** AI-generated material should not automatically become part of the authoritative reasoning record. The student must accept, reject, or modify it.
- **Authenticated actor identity.** Consequential actions should be associated with the student User ID or other authorized actor.
- **Time-ordered lifecycle history.** The system should preserve the order in which major reasoning steps, evidence additions, challenges, and decisions occurred.
- **Append-only revision history.** Material revisions should preserve prior states rather than silently overwrite them.
- **Evidence traceability.** Evidence used to support claims should be linked to an identifiable source or artifact appropriate to the assignment.
- **Stage integrity.** Earlier reasoning states should not simply be rewritten after the student knows the final answer; the lifecycle should preserve progression from initial position through challenge and closure.
- **Auditability of human–AI interaction.** The system should retain enough information for an instructor to distinguish student judgment from AI contribution at material decision points.
- **The governance floor is enforced; the analytical layer is adaptable.** Depending on age, subject and pedagogical purpose, components such as likelihood elicitation, or quantitative diagnostics may be simplified, omitted, or adapted. What should remain constant is the integrity of the reasoning record.

2.6 Why the Reasoning Record Is Difficult to Fabricate

A reasonable first question is whether a student could simply have AI generate the whole reasoning record — the counter-argument, the rejection rationale, the revision narrative — and submit an exemplary-looking lifecycle. The governance floor materially constrains this form of fabrication. It does not eliminate misconduct, but it changes the problem from reconstructing authorship after the fact to evaluating explicit, attributable actions within a governed record.

The reviewable objects are not prose the student writes about their reasoning. They are typed records produced by operations on the system. Origin is a structural field rather than a claim: material introduced by AI is recorded as AI-originated, so a student who asks AI for the strongest counter-explanation and admits it produces a record showing exactly that. AI contributions do not become authoritative on arrival; they remain pending until the student expressly accepts, rejects or modifies them, and that disposition carries an actor and a timestamp. Evidence is linked to identifiable sources, so the recency and quality attributes of the evidence set derive from the artifacts themselves rather than from what the student asserts. And because earlier states are preserved and the lifecycle progresses in order, an initial position cannot be quietly rewritten once the student knows where the analysis ended up.

The residual risk is out-of-band use: a student can generate material in a separate tool and enter it as their own. ARIP cannot detect that directly. What it does is change the nature of the problem. Entering AI-generated material under student origin is an affirmative misstatement on a specific, timestamped, attributable field — a conduct matter that institutional integrity processes already know how to handle — rather than an unresolvable inference about a finished document. It is also more laborious than doing the work, since a fabricated record must remain internally consistent across evidence, sources, sequence and disposition.

ARIP therefore does not claim to prove human cognition. It claims something narrower and more useful: that the student’s participation is recorded structurally rather than asserted, and that departing from it requires a deliberate and traceable act.

2.7 Illustrative Use Case: First-Year Business Course

Assignment: 'A Canadian manufacturer is considering automating part of its production process using AI-enabled robotics. Should management proceed with a $5 million investment?'

A conventional AI-enabled assignment might produce a polished recommendation immediately. In an ARIP environment, the student would instead be required to create and test the decision structure.

- Define a testable investment hypothesis.
- Create claims covering cost, productivity, workforce impact, implementation risk and strategic fit.
- Provide evidence for each claim and explain why it is credible.
- Ask AI to identify the strongest contradiction or alternative explanation.
- Decide which AI suggestions are useful and document rejections.
- Revise the hypothesis or evidence set where the contradiction is material.
- Close with a recommendation, uncertainty statement and explanation of how the reasoning changed.

What gets graded: The final recommendation can still be assessed for course knowledge. ARIP adds visibility into how the student arrived there — making reasoning, evidence discipline and human oversight part of the assessable work.

2.8 Institutional Value Proposition

- **For students:** Builds explicit practice in idea formation, evidence evaluation, contradiction handling and reasoned revision while learning to use AI as a cognitive tool rather than an answer generator.
- **For instructors:** Provides a richer assessment record than the final submission alone and supports feedback on reasoning process, not only subject-matter output.
- **For institutions:** Creates an auditable model for responsible AI use that is more constructive than blanket prohibition and more informative than post-hoc AI detection.
- **For academic integrity processes:** Produces attributable, timestamped records of what a student contributed and what AI contributed, addressing the evidentiary standard integrity bodies increasingly require and that AI-detection output alone does not meet.
- **For employers and society:** Potentially develops graduates who can challenge AI, explain decisions and exercise accountable judgment in AI-enabled workplaces.

3. Inclusive Reasoning and Neurodiverse Learning Styles

A central design for ARIP Education is to support different cognitive and reasoning styles without requiring the institution to label, diagnose or rank students by neurotype. Some learners may approach problems linearly; others may excel at divergent idea generation, intense focus, anomaly detection, pattern recognition, or making non-obvious connections across information.

A governed ARIP workflow gives those differences room to surface while preserving a common discipline for testing them. A student may begin with an unconventional connection or weak signal, but the idea must still be converted into an explicit proposition, tested against evidence, challenged by contradiction, revised where warranted, and brought to a reasoned closure.

This is potentially valuable for neurodivergent learners because educational systems do not always capture strengths that are expressed through non-linear reasoning, deep interest, unusual pattern detection, or alternative problem-solving approaches. ARIP could provide a structured environment in which those strengths become visible and assessable without relaxing evidentiary standards.

- Multiple paths to idea formation while preserving one auditable reasoning framework.
- Explicit capture of original student observations, hypotheses and pattern-based intuitions before AI expansion.
- A requirement to test unusual or non-linear insights against evidence and counter-explanations.
- Teacher visibility into how the learner developed, challenged and revised an idea rather than judging only the final presentation style.
- Potentially richer assessment for students whose reasoning strengths may not be fully reflected in conventional timed or output-only evaluation.

The objective is not to diagnose neurodivergence or imply that any cognitive trait belongs exclusively to a diagnostic category. Educational claims about benefits for autistic, ADHD, OCD or other neurodivergent learners would require dedicated research with those populations. The ARIP design principle is broader: ARIP allows diverse reasoning styles to be expressed, examined and assessed fairly while maintaining transparent standards for evidence, contradiction and human judgment.

4. Privacy, Governance and Safeguards

An education deployment could involve minors, sensitive assessment records, and potentially third-party AI services. Governance cannot be an afterthought. A deployment must address institutional privacy law, consent, retention, data residency, accessibility, accommodation, academic freedom and appeal mechanisms. Obligations vary materially by province, and are more demanding again where minors are involved, so a privacy and legal review scoped to the specific jurisdiction should precede any pilot rather than follow it.

- Clear disclosure of what student interaction data are captured and why.
- Role-based access separating student, instructor, administrator and research uses.
- Institution-controlled retention and deletion rules.
- No hidden automated grading or punitive action based solely on an unvalidated metric.
- Human review for consequential assessment decisions.
- Controls over which information can be sent to external AI providers.
- Accessibility and accommodation pathways that do not disadvantage students with different learning needs.
- Governed research protocols before using historical student data to calibrate scoring models.

4.1. Relevance to Canadian AI and Education Strategy

Canada's AI policy discussion increasingly includes adoption, productivity, responsible use and talent. Education introduces an additional strategic requirement: widespread AI use should not erode the human reasoning capacity on which advanced study, professional judgment and democratic institutions depend.

An ARIP Education model could contribute to a broader 'Human Reasoning and AI Literacy' pillar in which learners are not merely taught how to prompt AI, but how to question it, test evidence, identify contradiction, explain decisions and remain accountable for conclusions.

Responsible AI education should measure more than access to tools. It should create environments in which human reasoning remains visible, practiced and accountable.

5. Pilot Proposal

A prudent first deployment would be a bounded pilot in one course or faculty rather than an institution-wide rollout. The purpose would be to test usability, educator value, student behaviour and the feasibility of reasoning-provenance assessment.

Pilot Element	Illustrative Design
Scope	One course; 25–100 students; 2–4 governed assignments
Comparison	Where feasible, compare conventional assignment workflow with ARIP-supported workflow
Primary outcomes	Instructor usefulness, student engagement, reasoning visibility, completion burden, usability
Secondary outcomes	Quality of evidence handling, contradiction response, revision behaviour, academic-integrity observations
Validation	Compare any ARIP-generated indicators with independent educator ratings; do not treat indicators as validated until agreement and outcome evidence are established

Pilot Element	Illustrative Design
Governance	Institutional privacy, ethics and academic-integrity review before launch

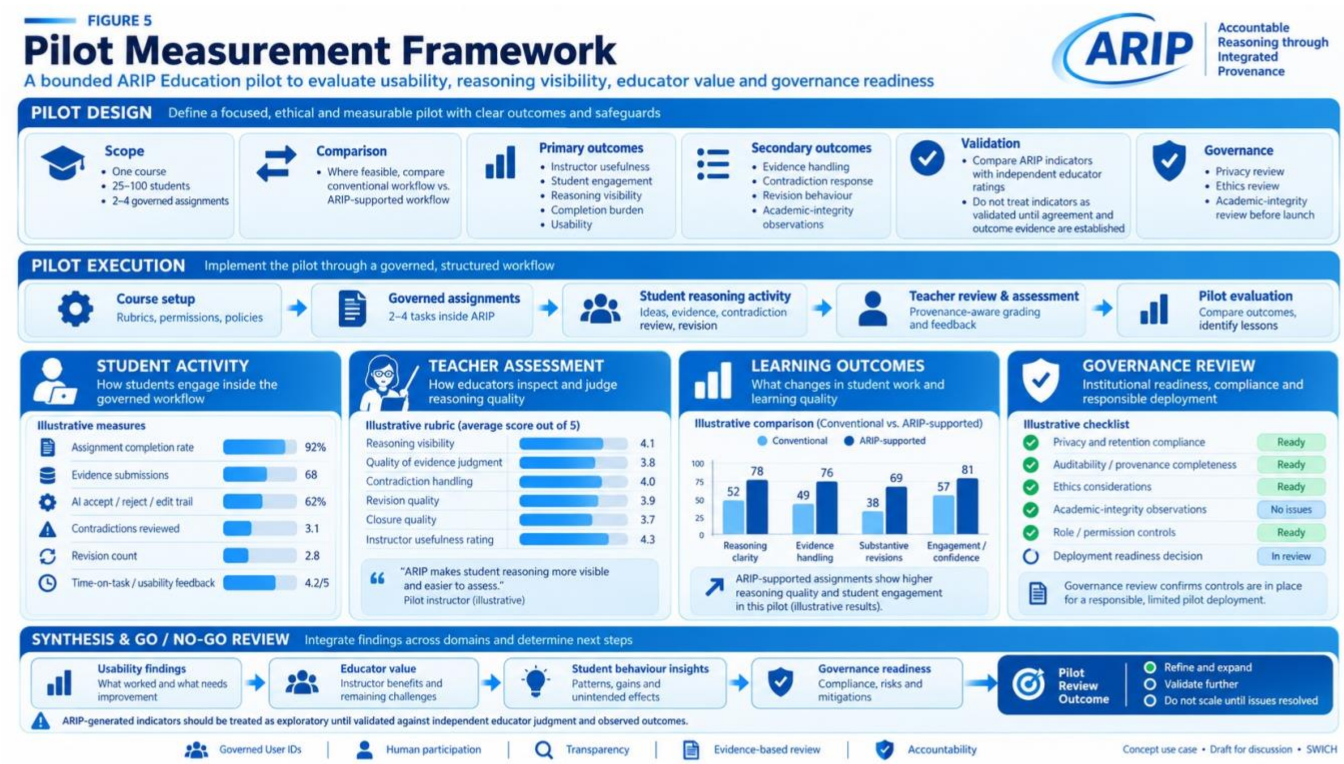

Figure 6. Pilot measurement framework

5.1 Commercial Deployment Model

ARIP Education would be built on ARIP Core rather than as a separate reasoning engine, but the reuse is uneven and the distinction matters for planning. The governance layer transfers directly: identity and role management, the artifact registry, bounded AI task packages, re-ingestion with preserved provenance, per-object User acceptance, the append-only event ledger, and enforced stage ordering. That layer is the substantial and difficult part of ARIP, and it supplies almost all of the minimum governance floor at §2.5. The analytical layer does not transfer in the same way. ARIP Core's quantitative machinery is calibrated for time-bounded speculative-event decisioning, and a pedagogically appropriate analytical layer for student work is new design rather than configuration. ARIP Education should therefore be planned as a governed reuse of ARIP Core's governance and provenance infrastructure, with a purpose-built educational analytical layer above it. Institution-specific deployments could combine tenant isolation, student and instructor identity management, assignment templates, teacher dashboards, governed AI-provider access and longitudinal ULA records.

Potential deployment units include:

- University course or department
- College program
- Secondary school
- School board
- Professional certification or continuing education
- Corporate graduate / leadership training

6. Conclusion

Generative AI makes it increasingly difficult to infer student thinking from a finished assignment. That does not require education to retreat from AI. It creates an opportunity to redesign the learning process around visible reasoning. ARIP Education treats the student's interaction with AI as part of the assessable learning journey. By preserving idea formation, evidence choices, contradiction, revision and closure, the platform could help educators evaluate whether AI is strengthening human judgment or replacing it.

ARIP does not ask only whether a student used AI.
It makes visible how the student reasoned with AI.

Appendix A — Illustrative Teacher Assessment Rubric Anchored to Governed ARIP Records

This rubric is illustrative only and would require educator design and validation before operational use. ARIP does not automatically determine whether a student has reasoned well. Instead, it provides a governed evidentiary record that allows the teacher to assess the student's intellectual participation, use of evidence, response to contradiction, treatment of AI contributions, and justification for closure.

The evidence available to the teacher depends on how the deployment is configured. Dimensions resting on origin attribution, acceptance decisions and lifecycle history are supported wherever the governance floor at §2.5 is preserved; dimensions drawing on likelihood elicitation or quantitative diagnostics assume a deployment that retains the fuller analytical layer. The objective is therefore not to reward particular system behaviours — such as rejecting more AI suggestions or changing an initial hypothesis, but to evaluate whether the student exercised appropriate judgment given the evidence available.

Dimension	Teacher judgment	Governed ARIP evidence available	Developing	Proficient	Advanced
Hypothesis formation	Did the student meaningfully define the problem rather than simply adopt an AI-generated framing?	Origin attribution; initial IDEA/DEFINE record; frozen prior state; hypothesis versions; assumptions and alternatives captured before evidence admission	Largely repeats the prompt or adopts an AI-proposed framing without meaningful refinement	Defines a clear, relevant and testable proposition with appropriate scope	Develops a precise, testable proposition; identifies material assumptions and credible alternatives; demonstrates meaningful student origination or refinement
Evidence judgment	Did the student select, evaluate and use evidence appropriately?	Evidence records; verifiability_enum; verifiability_scalar; source_authority_enum; reproducibility_enum; source artifacts and locators; inclusion/exclusion rationales; temporal fields	Uses evidence with limited attention to source quality, relevance or limitations	Selects relevant evidence and provides reasonable justification for credibility and use	Compares competing evidence, identifies limitations, distinguishes stronger from weaker sources, and uses evidence proportionately to its diagnostic value
Contradiction handling	Did the student seriously engage evidence or explanations that challenged the preferred position?	Contradiction objects; origin of alternative explanation; contradiction-side likelihood treatment; M-BCD first-pass/final states; accepted alternatives; subsequent hypothesis/PAM movement	Minimizes, ignores or dismisses material contradiction without adequate analysis	Addresses significant contradiction and explains why it does or does not alter the hypothesis	Actively engages the strongest counter-explanation, evaluates it fairly, and updates — or deliberately retains — the position in proportion to the evidence
AI oversight / intellectual control	Did the student exercise judgment over AI contributions rather than passively accepting them?	AI-origin typing; AI_ASSISTED records; acceptance/rejection/revision actions; actor and timestamp; rejection or modification rationale; pending/accepted state transitions	Frequently accepts AI contributions without meaningful scrutiny or justification	Questions, edits, rejects or accepts AI suggestions with reasonable explanation	Systematically evaluates AI contributions, identifies errors or weaknesses, improves useful suggestions, rejects inappropriate ones, and demonstrates clear intellectual control over the reasoning

Dimension	Teacher judgment	Governed ARIP evidence available	Developing	Proficient	Advanced
					process
Response to challenge	Did the student change the analytical position when warranted — or appropriately retain it when it survived challenge?	First-pass versus final likelihoods; append-only revision history; pam_movement_enum; delta_pam_support; sequential PAM ledger; changes to claims/evidence following contradiction	Shows little relationship between new evidence and subsequent reasoning; changes or retains position without adequate justification	Revises the position when material evidence warrants change, or provides defensible reasons for retaining it	Demonstrates disciplined stability or revision under challenge: analytical movement is proportionate to the strength of new evidence and the student clearly explains why the position changed or remained stable
Reasoning progression	Did the student's reasoning become more coherent and evidence-based as the lifecycle progressed?	Time-ordered lifecycle records; successive claim/evidence states; accepted revisions; M-BCD changes; ANALYZE run lineage; append-only stage history	Reasoning changes appear disconnected, retrospective or insufficiently supported	Shows a logical progression from initial position through evidence and analysis	Demonstrates a coherent, traceable progression in which major changes can be linked to specific evidence, contradiction or analytical findings
Closure and proportionality	Is the final conclusion justified by the governed analytical record?	Final PAM support state; unresolved contradiction; evidence record; closure rationale; uncertainty statement; CONTROL/CLOSURE record	Conclusion materially exceeds, understates or ignores the analytical record	Conclusion generally reflects the evidence and acknowledges important limitations	Conclusion is proportionate to the governed record, transparent about uncertainty and unresolved issues, and clearly explains why closure is justified
Reasoning provenance / participation	Does the record demonstrate substantive student participation throughout the lifecycle?	Actor identity; origin typing; acceptance decisions; time-ordered actions; student-authored versus AI-assisted objects; append-only audit history	Record shows limited meaningful student intervention or heavy reliance on externally generated reasoning	Record demonstrates recurring substantive student judgment across the lifecycle	Record demonstrates sustained, consequential student participation in defining, testing, challenging and closing the hypothesis

A1. Interpretation Principle

ARIP evidence should **inform teacher judgment, not replace it**. A high number of AI rejections does not automatically indicate strong reasoning. A student may correctly accept a strong AI contribution. Likewise, substantial revision is not inherently superior to stability. A well-formed initial hypothesis may survive serious contradiction and remain essentially unchanged.

The relevant educational question is therefore: **Did the student respond appropriately to the information, evidence and AI contributions encountered during the lifecycle?**

A2. Stability under challenge

ARIP Education explicitly recognizes two legitimate forms of strong performance:

- **Evidence-based revision:** the student changes a claim, hypothesis or conclusion because stronger evidence or contradiction warrants it.
- **Evidence-based stability:** the student seriously evaluates contradictory evidence but retains the original position because it remains better supported.

This distinction prevents the rubric from rewarding change for its own sake.

A3. Why this rubric is different from conventional AI assessment

A conventional rubric generally evaluates what the student submitted. This ARIP-supported rubric can evaluate **what the student did during the reasoning process**, because the relevant evidence is preserved structurally rather than reconstructed retrospectively from prose.

The teacher can therefore distinguish between:

- AI-generated content that the student passively accepted;
- AI-generated content that the student critically evaluated and adopted;
- student-originated reasoning that survived challenge;
- student reasoning that was appropriately revised;
- conclusions that are consistent — or inconsistent — with the analytical record.

The system does not merely show who produced the words. It provides evidence of who exercised the judgment.

Appendix B — Competitive Landscape for ARIP Education

This appendix summarizes an initial market scan of systems that overlap with the proposed ARIP Education model. The review identified meaningful competition in AI-assisted writing, authorship provenance, Socratic tutoring, student-process visibility, and institution-controlled AI learning. It did not identify, in this scan, a widely deployed product that clearly combines those functions inside an end-to-end governed reasoning lifecycle comparable to the proposed ARIP model.

B.1 Competitive Categories

ARIP Education would compete across several adjacent product categories rather than against a single established category. This matters strategically: if ARIP is positioned only as an AI tutor, authorship detector, or writing-transparency tool, it enters markets where mature competitors already have strong institutional distribution.

Category	Representative Systems	Primary Question	ARIP Distinction
Authorship / integrity	Turnitin Clarity; Grammarly Authorship	How was this work produced, and what role did AI play?	Govern the reasoning process, not only textual authorship or drafting provenance.
AI tutoring / personalized learning	Khanmigo; SchoolAI; similar tutor platforms	How can AI help this student learn?	Make the student's evidence, contradiction handling, revision, and closure explicitly assessable.
Socratic / critical-thinking systems	Socrado; SchoolAI; Materact	How can AI prompt the student to think rather than simply answer?	Persist a governed reasoning object across the entire assignment lifecycle to closure.
Institution-controlled AI learning	Ontario Tech AI Learning Agent and comparable institutional agents	How can institutions provide trusted, accountable AI support?	Add cross-domain hypothesis/evidence governance, provenance, teacher review, and longitudinal learning.

B.2 Closest Competitors and Adjacent Systems

Turnitin Clarity

Closest commercial overlap in process transparency. Clarity provides a dedicated writing environment in which the student's drafting process is visible to instructors as 'proof of process'; instructors can review writing activity, timeline/playback, AI-chat activity when enabled, and assignment-level AI permissions. Its center of gravity remains writing, originality, and academic integrity.

ARIP positioning: ARIP differentiates by governing the intellectual path from idea to hypothesis with evidence, contradiction, revision, and closure across disciplines, not merely the production history of a written artifact.

Grammarly Authorship

Strong overlap on provenance. Authorship tracks whether text is student-written, copied, or AI-generated and supports full assignment replay. This gives instructors transparency into composition and source use.

ARIP positioning: ARIP avoids competing on a simplistic 'human percentage.' Its stronger claim is intellectual contribution: whether the learner exercised judgment over claims, evidence, counter-explanations, AI suggestions, and final conclusions.

SchoolAI Spaces

SchoolAI combines personalized AI interaction with educator visibility. Teachers can see student conversations in real time, and the company positions Spaces around helping students do the thinking rather than merely receiving answers.

ARIP positioning: ARIP differentiates through explicit governed stages, structured evidence objects, required contradiction testing, formal closure, and longitudinal reasoning records rather than primarily conversational personalization.

Socrado

Socrado is philosophically close to the ARIP Education concept. It uses Socratic dialogue to develop critical thinking, supports instructor dashboards, learning goals, and progress tracking, and is designed specifically for educational use rather than generic chat.

ARIP positioning: ARIP's distinction is a persistent end-to-end reasoning record that can be audited and assessed as a structured artifact, not only a sequence of instructional dialogue.

Materact

Materact is especially relevant as evidence that educators value process-level visibility. In mathematics, it analyzes lines of student working, uses Socratic AI support, lets teachers see conversations, and maps reasoning steps to curriculum objectives.

ARIP positioning: Its domain specificity creates whitespace for a cross-domain ARIP model that can structure reasoning in business, science, humanities, policy, research, and other evidence-based assignments.

Ontario Tech AI Learning Agent

Ontario Tech's 2026 pilot is strategically important in the Canadian context. The university describes an in-house, course-based AI learning system designed around trust, accountability, academic integrity, and support for human judgment across undergraduate and graduate courses.

ARIP positioning: This is treated as both an adjacent institutional model and market validation. ARIP can demonstrate why its formal reasoning lifecycle, auditability, and cross-assignment learning add value beyond institution-grounded AI assistance.

Khanmigo / broader AI tutor market

Khanmigo and the wider AI-tutoring category compete for student interaction and teacher-supported learning. These platforms benefit from broad instructional ecosystems and established user bases.

ARIP positioning: ARIP is not a general tutor. Its category is reasoning governance: making the student's reasoning process visible, testable, attributable, and reviewable.

B.3 Potential Competitive Whitespace

Across the systems reviewed, individual elements of the ARIP Education concept are already well represented in the market. The opportunity is therefore not that ARIP has invented AI-assisted learning, process transparency, or Socratic tutoring. The potential whitespace lies in the combination and governance of those capabilities.

- Persistent student identity linked to an institution-controlled reasoning record.
- Idea or hypothesis formation as a first-class activity rather than beginning only with a fixed answer prompt.
- Explicit claims, evidence, source quality, and counter-explanations.
- Required contradiction testing before closure.
- Recorded student acceptance, rejection, or modification of AI contributions.
- Reasoning revision across a governed lifecycle.
- Teacher assessment of both subject output and intellectual process.
- Formal closure that records why the learner reached the conclusion and what uncertainty remains.
- Longitudinal cross-assignment learning through a ULA-style observational layer.
- Cross-domain applicability rather than restriction to writing or mathematics.

Competitive claim guardrail: The current market scan suggests that ARIP appears differentiated by its proposed combination of governed reasoning stages. It does not suggest a universal claim that 'no competitor does this.' Competitive capabilities evolve quickly, so the landscape must be refreshed frequently.

B.4 Strategic Positioning

The initial positioning avoids defining ARIP Education as an AI-detection or tutoring product. A stronger category statement is:

ARIP Education is a governed human–AI reasoning environment that allows institutions to observe, structure, and assess how learners progress from an idea or assigned problem through evidence, contradiction, revision, and conclusion.

A concise competitive distinction is:

Authorship tools show how a document was produced. AI tutors help a student learn. Socratic systems encourage the student to think. ARIP is designed to make the student's governed reasoning journey itself an auditable and assessable object.

B.5 Implications for the ARIP IDEA Lifecycle

The existence of strong adjacent products becomes evidence inside the ARIP Education lifecycle rather than treated only as a marketing consideration. The test is whether the proposed differentiation is sufficiently material to justify a separate ARIP Education module and whether institutions would value a full reasoning-governance layer beyond existing writing, tutoring, and academic-integrity tools.

- Can ARIP demonstrate materially better visibility into intellectual contribution than authorship-provenance tools?
- Does a formal IDEA-to-CLOSURE workflow improve reasoning or merely increase student and teacher workload?
- Which subjects benefit most from hypothesis/evidence/contradiction governance?
- Will educators pay for a separate reasoning layer if their LMS, Turnitin, or AI tutor already provides process visibility?
- Can teacher review remain practical at classroom scale?
- Can ARIP integrate with existing LMS and integrity systems rather than requiring institutions to replace them?
- Do longitudinal ARIP records provide educational value that existing dashboards do not?

Sources

- **Cherki El Idrissi, S. (2025).** *Enhancing or Hindering? The Influence of Generative AI on Critical Thinking and Collaborative Learning.* AMCIS 2025 Proceedings. University of Toronto Mississauga.
- **Philion, E., Jeanneret, C., & Gales, J. (2026).** *Understanding the Use of Generative Artificial Intelligence (GenAI) by Students in an Undergraduate Civil Engineering Capstone Course.* Proceedings of the Canadian Engineering Education Association. York University.
- **Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025).** *Generative AI without guardrails can harm learning: Evidence from high school mathematics.* Proceedings of the National Academy of Sciences, 122(26), e2422633122. https://doi.org/10.1073/pnas.2422633122
- **Lehmann, M., Cornelius, P. B., & Sting, F. J. (2025).** *AI Meets the Classroom: When Do Large Language Models Harm Learning?* SSRN Working Paper 4941259. University of Cologne and Rotterdam School of Management, Erasmus University. https://doi.org/10.2139/ssrn.4941259
- **MIT Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training (2026).** *Report of the Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training.* Massachusetts Institute of Technology, August 13, 2026. https://aiandeducation.mit.edu/report/

Appendix C — Questions for the ARIP Development Partner

This appendix is the purpose of circulating this paper at this stage. It is not a scope document, a requirements specification or a request for an estimate. ARIP Education is a concept under evaluation, and no commitment is sought or implied.

What is sought is architectural input. Several decisions being taken during the Phase 1 ARIP-SIM build are inexpensive to accommodate now and expensive to retrofit later. If ARIP Education proceeds as the first commercial deployment after ARIP-SIM, the questions below are the ones whose answers would change what Phase 1 should leave open.

C.1 Governance floor and reuse

Floor coverage. Of the minimum governance floor at §2.5 — origin attribution, explicit student disposition of AI contributions, authenticated actor identity, time-ordered lifecycle history, append-only revision, evidence traceability, stage integrity and auditability of human–AI interaction — which requirements does the Phase 1 architecture already satisfy, which would be satisfied with minor extension, and which would need to be designed in now rather than added later?

Reuse boundary. §5.1 assumes the governance and provenance layer is reusable and that the educational analytical layer is new design. Does that division match how the Phase 1 codebase is actually structured, and is the boundary between governance and analytics clean enough to support a second analytical layer without forking Core?

Per-object acceptance. Technical Specification §5.15.5 already specifies one reviewable candidate object per parsed output, each with its own acceptance record. How closely does that mechanism match what an educational review interface would need, and what would have to change?

C.2 Deployment architecture

Multi-tenancy. Does the Phase 1 design assume a single tenant? What would institutional multi-tenancy imply for data separation, configuration and operational cost, and is it cheaper to establish the tenancy model now than to introduce it later?

Identity federation. Institutions will expect single sign-on against their own directory, with instructor, student and administrator roles derived from institutional records. How far does the Phase 1 identity model extend toward that?

AI provider abstraction. Institutions increasingly route AI access through their own governed gateway rather than accepting a vendor's provider relationship, and this also bears on Canadian data residency obligations. Does the External Chat Mode and API Mode split accommodate an institution-supplied gateway cleanly, and what would be required to make the provider fully substitutable?

LMS integration. Assignment distribution, roster synchronisation and grade return would normally run through the institution's learning management system. Is an LTI-style integration the right route, and what does it assume about the Phase 1 API surface?

C.3 Review at scale

Instructor workload. The binding constraint on adoption is instructor review time, not reasoning quality. A realistic target is a median of under ten minutes per student per assignment. What would the review interface have to do — in triage, summarisation and exception surfacing — to make that achievable against a governed lifecycle record?

Cohort view. What is required to present an instructor with a cohort-level view that identifies the small number of lifecycles warranting close attention, without implying automated assessment of reasoning quality?

C.4 Sequencing and dependencies

Phase relationship. What would need to be true at ARIP-SIM completion for an Education pilot to begin without rework? Which elements could run in parallel with Phase 1, and which strictly follow it?

Deferred capabilities. The longitudinal learning record at §2.4 depends on ULA, which is presently deferred. Does that place the longitudinal capability firmly in a later phase, and does anything in Phase 1 need to anticipate it?

Cross-domain support. ARIP-SIM is currently the only supported domain, and cross-domain enumeration support is deferred. What is the practical effort in standing up a second domain module, and does that work belong in Phase 1 or later?

C.5 What would change your view

Risks and alternatives. Where does this concept look weakest from an engineering standpoint? If ARIP Education is not the right second deployment, what would be, and what in this paper would need to be true for that assessment to change?