

AI-Resilient Learning Initiative

A Framework for Assessment in the Age of Generative AI

Contents

Glossary	3
1 Program Overview	4
1.1 What This Program Does	4
2 The Three AI-Resilience Pillars	4
2.1 Scope and Limits	5
2.2 Outcomes We Want to See.....	5
2.3 Student Involvement.....	6
What student partners do:.....	6
Why student involvement matters:.....	6
2.4 AI-Resilience Process Flow.....	6
1.4.1 Faculty Track:.....	6
2.4.2 Student track	7
3 The AI-Resilient Assessment Framework	7
3.1 How Bloom's Taxonomy Connects to AI Resilience	8
3.1.1 Bloom's as a Diagnostic Lens	8
3.1.2 Why Lower-Level Tasks Are Often More Exposed	8
3.1.3 The Problem with Product-Only Assessment	9
3.1.4 Higher-Level Tasks Are Not Automatically Secure	9
3.1.5 Why Analyse (C4) Matters in This Framework	9
3.1.6 How the Three AI Resilience Pillars Respond	9
3.2 AI Vulnerability Spectrum.....	9
3.3 From Vulnerability to Redesign	10
3.3.1 Identifying the Assessment Fracture	10
3.3.2 Diagnosing the Main Source of Vulnerability.....	10
3.3.3 Selecting the Primary Design Response	10
3.3.4 From Primary Response to Full Redesign.....	10
4 AI-Vulnerable vs. AI-Resilient Task Design.....	10
5 Pillar Integration — A Non-Chronological Framework.....	11

Glossary

Terms used consistently across the AI-Resilience framework.

Term	Definition
AI-Resilient Assessment	A task designed so it cannot be completed by substituting AI output for genuine student cognitive engagement; requiring process evidence, real-world grounding, and structural conceptual understanding.
AI-Vulnerable Task	A task AI can complete easily, demonstrating a high standard of surface fluency without any genuine understanding, e.g. 'Explain the concept of X' or 'Give examples of Y'.
Conceptual Knowledge	In Bloom's knowledge dimension, conceptual knowledge is the understanding of how ideas, categories, principles, and models relate to one another within a larger structure. It goes beyond remembering information by requiring students to grasp patterns, connections, and the underlying organisation of a concept.
Process>Product (P1)	Visible evidence of the cognitive process behind the student's output; drafts, concept maps, reflective annotations; which AI-generated output has no equivalent of.
Contextualization (P2)	Grounding a task in a specific, personally observed real-world context, the student can name, date, and describe in detail; details which are inaccessible to AI.
Conceptual Depth (P3)	Identifying a concept of essential structural components versus incidental ones; a discrimination AI cannot authentically demonstrate.

1 Program Overview

1.1 What This Program Does

The AI-Resilient Learning Initiative (AI-RLI) is a structured intervention in the Student Pedagogy Partnership Program by the Learning Innovation Center that is designed to protect the integrity of assessment in the age of generative AI. Rather than prohibiting AI, the program redesigns assessment to be resistant by AI-generated output; requiring evidence of genuine student thinking; evidence which is difficult to fabricate convincingly without genuine engagement.

AI-RLI operates through a **co-development clinic model** in which AI-Resilient Learning Lead work directly with faculty partners and students to evaluate existing assessments and redesign tasks using a structured three-pillar framework. (i) Process>Product, (ii) Contextualisation, and (iii) Conceptual Depth. This initiative will be run in accordance with the Student Pedagogy Partnership (SPP) and piloted in Spring 2026, establishing a model where teachers and students serve as "co-pilots" to ensure assessments are both authentic and effective in accordance with the three-pillar framework of AI-resilience. Student partners extend their help, in building peer-level awareness of AI-resilience as a learning practice, in parallel to their faculty partners.

2 The Three AI-Resilience Pillars

Each pillar targets a different dimension of Learner -AI collaboration. The application of these pillars alone is insufficient; together they create an assessment that AI output has a hard time substituting for, retaining the learner's agency of critical thinking and learning command

Pillar	What it does	Core design move
P1 Process > Product	Requires students to produce evidence of thinking showing how they approach the assessment against the concept and its application in-context. Since, AI has the ability to produce polished outputs with no cognitive arc, no confusion, revision, or moment of insight, this pillar allows to capture that cognitive work necessary for assessment against a concept that was learned.	Require drafts, concept maps, plans, or reflective annotations showing the thinking behind the output. Creates evidence that AI might not be able to reverse engineer.
P2 Contextualization	Requires students to present a real, personally observed context and apply the concept to it. AI has no access to a student's specific lived-experience and can only produce generic examples.	Require a specific, traceable, verifiable, course-relevant context. Students must document real execution and reflect on how the concept operates against real-world constraints.

<p style="text-align: center;">P3 Conceptual Depth</p>	<p>Tests whether students have internalized the essential structure of a concept, the minimal set of components that are non-negotiable to constitute that concept. AI can describe concepts fluently but cannot show which elements are loadbearing.</p>	<p>Require students to distil a concept to its essential components, justify why each is non-negotiable, identify components often mistaken as essential, and transfer the structure to a novel context. Students will use AI Critique as a Conceptual Depth Strategy</p> <p>Essential structure of AI Critique Students will have to ask: Where could the AI response be better? What did it oversimplify, miss, or get wrong? They submit a paragraph critiquing the AI response against their own structural analysis.</p>
--	---	---

2.1 Scope and Limits

AI-RLI does **not** claim to make assessment "AI-proof," nor does it assume that all AI use is unacceptable; rather, it helps faculty decide where AI use may be acceptable, where it must be disclosed, and where it undermines the validity of the task. In this sense, the initiative is concerned with preserving the integrity of learning outcomes, neither with enforcing a universal ban on AI tools nor strict punitive measures against the ethical use of AI

AI-RLI supports the redesign of assessments where the final product alone no longer provides reliable evidence of student learning, instead it mandates the distillation of concept being assessed into its essential components, conceptual application into a context and evidence of thinking showing how they approached the assessment

Because disciplines differ in method, evidence, and assessment culture, AI-RLI is intended to be applied **adaptively rather than mechanically**. The same three-pillar logic remains constant across disciplines, but the form of evidence, the type of context used, and the balance between process, application, and abstraction may vary between subjects. The framework therefore provides a shared structure for redesign while allowing discipline-specific interpretation.

2.2 Expected Outcomes through the application of AI-resilience Pillars

Faculty Capability

Faculty independently design AI-resilient assessments using the three-pillar framework.

Student Literacy

Students develop AI-resilience as a learning practice. They are also able to recognise when AI supports

Institutional Evidence

Provide findings on AI usage patterns, assessment fractures, faculty development and impact.

Design tasks that require clearer evidence of student reasoning, judgement, and conceptual understanding

learning and when it bypasses the intended learning process.

Identify which redesign approaches are most workable across disciplines.

2.3 Student Involvement

AI-RLI is not a faculty-only initiative. Student partners are embedded in the process from the start, not brought in as passive recipients at the end.

What student partners do:

1. Support peer-level awareness of AI-resilience practices; reaching students in ways that faculty cannot.
2. Participate in pre-and post-intervention surveys to document shifts in their own AI usage and understanding,
3. Provide feedback on redesigned assessments from a student perspective before they are formally deployed,
4. Act as a bridge between the LIC and the wider student body, flagging where AI-resilience feels unclear or unreasonable from a learner's point of view.

Why student involvement matters:

Faculty redesign the task, students are the ones who must engage with it and they are also the ones most aware of how AI is being used by their peers. Involving students in the process builds buy-in, surfaces blind spots in the redesign, and develops a cohort of students who understand AI-resilience as a learning practice rather than a restriction.

Student partners are recruited through the SPP and work alongside faculty partners throughout. Their observations and survey data form part of the final findings report of the pilot initiative.

2.4 AI-Resilience Process Flow

1.4.1 Faculty Track:

The faculty track runs left to right across four colour-coded phases (Please refer to the AI-RLI Process flow below):

1. **Preparation** (Steps 1–3): where faculty are oriented and course outlines are sourced and evaluated;
 2. **Design** (Steps 4–5): where individual feedback meetings and weekly co-development clinic sessions take place;
 3. **Delivery** (Steps 6–7): where students are equipped with AI-resilience practices and the redesigned assessment is rolled out; and
 4. **Evaluation** (Steps 8–10): where results are discussed, a post-intervention survey is conducted, and final findings are analysed.
-

Running beneath Steps 2–6 concurrently: faculty conduct an AI-allowed assessment during the design phase. This is not a sequential step; it runs alongside the clinic sessions and surfaces real AI usage patterns and assessment fractures that directly inform what gets redesigned in Step 7.

2.4.2 Student track

The students track shows five touchpoints where students actively participate:

1. **Orientation:** An orientation briefing at Step 1 (including the pre-survey baseline);
2. **AI Allowed task:** Engagement with the AI-allowed task during the parallel step;
3. **Peer Awareness:** A peer awareness activity at Step 6 where student partners brief their peers;
4. **AI-Resilient Task:** Completion of the AI-resilient assessment at Step 7; and a post-survey at Step 9 that documents shift in AI usage and understanding.

The process flow below maps the AI-Resilience Learning Initiative programme:

AI-RLI Process Flow -SPP Spring 2026									
FACULTY									
PREPARATION			DESIGN		DELIVERY		EVALUATION		
1 Orientation	2 Course outline + assessment	3 Evaluate scope	4 Feedback meetings	5 Co-develop clinic	6 Equip students	7 Roll out assessment	8 Discuss results	9 Post-survey	10 Final report
STUDENTS									
Orientation Briefed on AI-RLI Pre-survey baseline			AI-allowed task Parallel: fractures surface		Peer awareness Student partners brief peers	AI-resilient task Engage with redesigned assessment		Post-survey Fill Post AI-RLI Assessment survey	

3 The AI-Resilient Assessment Framework

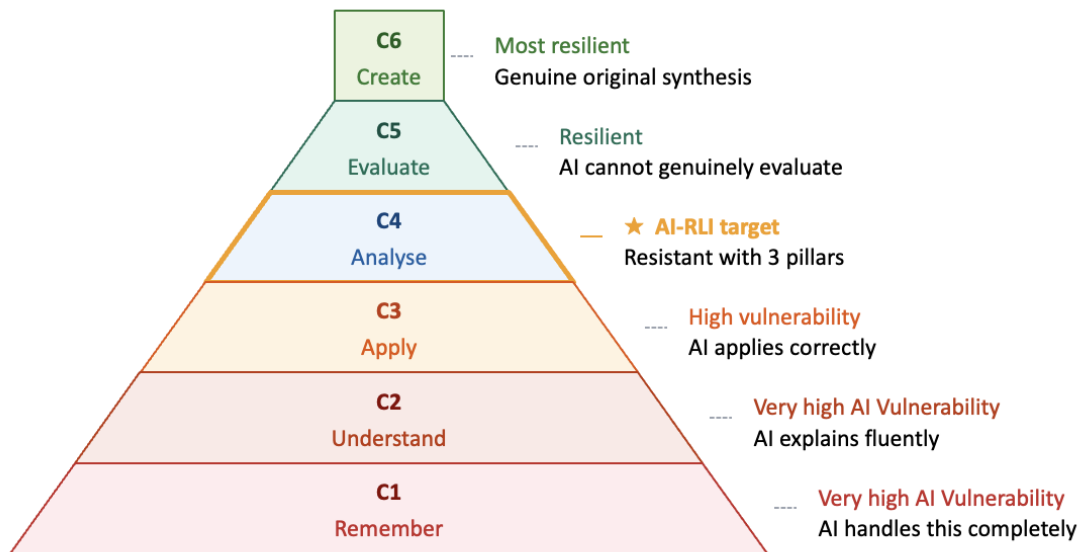
Bloom’s Taxonomy is a framework for classifying the type of cognitive work an assessment asks students to demonstrate. In this framework, the levels are referred to as **C1 Remember, C2 Understand, C3 Apply, C4 Analyse, C5 Evaluate, and C6 Create**. These levels move from recalling or explaining knowledge to using it, examining it, judging it, and producing something new from it.

Bloom's is important in this AI-resilient assessment framework because it helps faculty identify **what kind of thinking a task is designed to assess**. This matters in an AI-mediated context because different kinds of cognitive demand are not equally exposed when a task relies mainly on a polished final product. Bloom's therefore gives the framework a clear and shared way of locating the assessment demand, identifying where AI-supported substitution may be more likely, and guiding redesign so that student reasoning, judgement, and conceptual understanding are made more visible.

The diagram below illustrates blooms and its levels in correspondence with AI Resilience:

Bloom's Cognitive Levels

Higher levels resist AI substitution · Lower levels are fully AI-replaceable



3.1 How Bloom's Taxonomy Connects to AI Resilience

3.1.1 Bloom's as a Diagnostic Lens

Bloom's Taxonomy is used in this framework as a diagnostic lens for identifying where an assessment may be more vulnerable to AI-supported substitution. It helps explain why some tasks are easier for generative AI to complete with surface fluency than others, but it is not, on its own, a sufficient basis for designing AI-resilient assessment. The resilience of a task depends not only on its cognitive level, but also on how learning is evidenced, how specific the context is, what conditions the task is completed under, and what the marking criteria reward.

3.1.2 Why Lower-Level Tasks Are Often More Exposed

Tasks focused primarily on Remember (C1), Understand (C2), and Apply (C3) often ask students to recall information, explain concepts, or use known procedures in familiar ways. These remain important for learning, but in assessment they are often more exposed to AI substitution when the task asks only for a polished final product. In such cases, generative AI can often produce responses that appear coherent, accurate, and academically credible without offering reliable evidence of the student’s own reasoning.

3.1.3 The Problem with Product-Only Assessment

The central issue is not that lower-level cognition lacks value, nor that AI always performs these tasks well. The issue is that many C1–C3 assessments, when designed as product-only tasks, provide limited visibility into how the student arrived at the answer. Where there is no requirement for process evidence, contextual judgement, or structural explanation, it becomes harder to distinguish genuine student understanding from AI-assisted completion.

3.1.4 Higher-Level Tasks Are Not Automatically Secure

This framework does not treat Bloom’s as a simple threshold model in which lower levels are unusable and higher levels are automatically secure. A task at Analyse (C4), Evaluate (C5), or Create (C6) can still be highly vulnerable if it is broad, generic, weakly evidenced, or rewarded mainly for fluent presentation. Equally, a task that includes C1–C3 learning may become more resilient when it is anchored in traceable context, supported by visible process evidence, and assessed through criteria that value judgement, selection, and justification rather than output alone.

3.1.5 Why Analyse (C4) Matters in This Framework

Within AI-RLI, Analyse (C4) is often a useful entry point for redesign because it asks students to distinguish essential from incidental features within a concept and justify those distinctions. This makes it easier to design tasks that reveal conceptual understanding rather than summary alone. However, C4 is not a guarantee of resilience. What matters is whether the assessment requires students to demonstrate analysis through evidence, context, and defensible judgement.

3.1.6 How the Three AI Resilience Pillars Respond

Bloom helps identify where vulnerability may lie; the Three Pillars provide the design response. Process > Product (P1) makes the student’s reasoning and development visible. Contextualization (P2) grounds the task in a specific and course-relevant context rather than a generic response space. Conceptual Depth (P3) asks students to identify what is essential within a concept, justify their choices, and demonstrate structural understanding. In this way, Bloom’s supports diagnosis, while the pillars support redesign.

3.2 AI Vulnerability Spectrum

Bloom's Cognitive Dimension -AI Vulnerability Spectrum					
◀ High vulnerability		Cognitive level		AI Resilience ▶	
C1 Remember	C2 Understand	C3 Apply	C4 Analyse	C5 Evaluate	C6 Create

			◀ Framework		
Very high	Very high	High	Threshold	Resilient	Most res.
AI handles this completely	AI explains fluently & accurately	AI applies procedures & formulas	AI-resilient with the 3 pillars	Requires defended personal judgment	Requires genuine original synthesis

▲ AI-RLI
operates here

3.3 From Vulnerability to AI-resilient Redesign

3.3.1 Identifying the Assessment Fracture

The first step in redesign is to identify the point at which AI-supported output can substitute for the learning the task is intended to assess. In AI-RLI, this point is referred to as the **assessment fracture**: the place where a polished response can be produced without giving reliable evidence of the student’s own reasoning, judgement, or conceptual understanding. Identifying the fracture helps clarify what the redesign must protect.

3.3.2 Diagnosing the Main Source of Vulnerability

Once the fracture has been identified, the task should be diagnosed more precisely. In most cases, vulnerability arises because one or more of the following is missing: visible evidence of process, a specific and traceable context, or conceptual depth beyond surface explanation. This diagnosis matters because different weaknesses require different redesign responses.

3.3.3 Selecting the Primary Design Response

The Three Pillars function as design responses to different kinds of vulnerability. Where the task relies too heavily on a polished final product, **Process > Product (P1)** should be strengthened. Where the task invites generic or decontextualized answers, **Contextualization (P2)** should be strengthened. Where the task rewards fluent explanation without testing structural understanding, **Conceptual Depth (P3)** should be strengthened. In practice, robust redesign usually requires all three pillars, but identifying the primary weakness gives faculty a clearer point of entry.

3.3.4 From Primary Response to Full Redesign

Redesign does not end with the selection of one pillar. A task becomes more resilient when all three pillars are integrated in a way that fits the discipline, the learning outcomes, and the assessment conditions. The purpose of the diagnosis is therefore not to isolate a single solution, but to identify where redesign should begin and how the full task should be strengthened.

4 AI-Vulnerable vs. AI-Resilient Task Design

AI-Vulnerable task	AI-Resilient redesign
<i>"Analyse [concept] and explain how it applies to [topic] or give examples."</i>	<ul style="list-style-type: none"> Identify essential vs. irrelevant components with structural justification. Critique the AI response. (P3)

AI-Vulnerable task	AI-Resilient redesign
AI completes this with high surface fluency. No pillar protects it.	<ul style="list-style-type: none"> • Describe a personally observed real-world context where the concept operated. Name it, date it. (P2) • Submit evidence the context was real. Explain your process. Describe challenges and the way they were resolved. (P1)

5 Pillar Integration -A Non-Chronological Framework

The three pillars: (P1) Process>Product, (P2) Contextualization, and (P3) Conceptual Depth, are not steps in a sequence. They are three **interdependent design lenses** that must all be present in the final assessment task, but there is no required order for integrating them. A faculty member can enter from any of the three, depending on the concept, the discipline, and what naturally comes to mind first.

They might start with a vivid real-world context students have access to (enter via P2), work out which essential components would be visible in that context (P3), then decide what evidence would prove it was real (P1). Or start with what evidence is realistic for students to produce (P1) and build backwards. All three routes produce the same task structure.