CALIFORNIA LUTHERAN UNIVERSITY

Teaching with Artificial Intelligence

MODULE 2

Friction, Productive Struggle, and Learning Design

Faculty Development Reading

Much of the time, the struggle is not in the way of the learning. The struggle is the learning.

— the course stance, in one sentence

June 2026

Module Purpose

Module 1 asked you to name what AI must not damage in your discipline. Module 2 supplies the learning-science spine that explains *why* certain damage happens — and why it happens unevenly. The central insight is uncomfortable in its simplicity: most of us have spent our careers designing assignments around the assumption that if a student turns in the work, the work gives us evidence of their thinking. Generative AI complicates that assumption, and it complicates it most for the students who have the least expertise to fall back on (Module 2 lectures).

This module's portfolio artifact is the **Friction Audit**: a one-page diagnostic of a single assignment that names the task, the intended learning, the essential friction, the AI bypass risk, one **Cognitive Sanctuary**, and one **Access Repair**. The major decision you are preparing to make is a sorting decision: *in this assignment, which difficulty is the learning, and which difficulty is merely in the way?* Both are friction. Only one is pedagogically valuable — and AI is excellent at removing both.

Learning Objectives

After completing this reading and the Module 2 activities, you will be able to:

1. Explain, in language a non-pedagogy colleague would understand, how productive struggle, retrieval, spacing, interleaving, and elaboration contribute to durable learning.
2. Identify a Safety Gap in one course artifact and explain why novices are more vulnerable to it than experts.
3. Diagnose the specific cognitive operation students must perform in one assignment before AI assistance is appropriate — using the course's six-operation menu.
4. Distinguish desirable difficulty from needless friction in that assignment, with at least one named example of each.
5. Design one Cognitive Sanctuary and one Access Repair, specifying what students do, what they submit, and how long it takes.

Before You Read: A 60-Second Retrieval Primer

Before the learning science, retrieve your own experience. Sixty seconds, no looking anything up.

> **YOUR 60-SECOND RESPONSE — write before you read**
>
> **Name one thing your students consistently struggle with — and your best current guess about why that struggle is happening.** One or two sentences. At the end of this reading you will sort that struggle into one of two very different categories.

Core Reading

The assumption AI breaks

Faculty assess what students submit; we infer what students know. The product is visible; the capability is invisible. For most of the history of higher education, that inference was reasonable, because producing the visible product required exercising the invisible capability. A coherent essay requires organizing an argument. Working code required understanding the algorithm. A clean lab report requires making sense of the data. Generative AI severs that linkage: the visible product can now be produced without the invisible capability being exercised at all. The product still exists. What it *proves* has changed.

The Safety Gap

The course names this severed linkage the **Safety Gap**: the distance between what a student can produce with AI assistance and what they could independently do without it. Two cautions keep the concept honest. First, a wide gap is not automatically a misconduct problem — it may reflect a reasonable scaffold, an accessibility support, or exactly the kind of assisted performance you intended. Second, the gap is dangerous chiefly when it is *invisible*: when neither you nor the student can say where assisted performance ends and independent capability begins. As the Module 2 lecture puts it, the Safety Gap is a visibility problem before it is a policy problem.

The gap looks different across disciplines but is structurally identical: a polished essay without argument ownership; correct code the student cannot debug; a smooth literature review the student cannot defend in conversation; a clean lab report the student cannot extend to a similar experiment. None of these examples begins with an accusation. Each begins with a quieter teaching question — *what does the submitted product still prove?* — and each points to the same design response: not detection, but a visibility plan. Somewhere in the course, the invisible capability has to become visible: through process, conversation, annotation, low-stakes quizzing, or brief oral explanation. Module 2 calls those moments Cognitive Sanctuaries, and you will design one.

Why novices are most exposed

Two pieces of cognitive architecture explain why the Safety Gap does not threaten all learners equally. **Working memory** is the small mental workspace where new information is processed — and it is severely limited (Sweller, 1988). **Schemas** are the organized long-term knowledge structures that let experts bypass that limit: a chess master sees one meaningful position where a novice sees thirty-two pieces. Schemas are what your courses exist to build, and they are built through effortful processing — retrieval, error, correction, elaboration — not through reading polished examples.

This is why the same AI interaction can help an expert and harm a novice *on the same task*. An expert reading AI-generated feedback has schemas against which to test it; she can spot the plausible-but-wrong suggestion and discard it. A novice has no internal model to test against, so fluent output simply reads as correct. The expert uses AI as an accelerator; the novice risks using it as a substitute for the very processing that would have built the schema. Mollick (2024, ch. 7) makes the same point from the

tutoring side: AI's educational promise depends on keeping the student doing the cognitive work, which is precisely what default chatbot behavior — give the answer, politely — does not do. The teaching implication is blunt: the earlier in your course, and the newer the material, the more protective your design needs to be.

What makes learning durable

If schemas are built through effortful processing, which efforts matter most? The learning-science literature converges on a family of **desirable difficulties** (Bjork & Bjork, 2011): conditions that feel harder in the moment and produce more durable learning over time. Four matter most for this course. **Retrieval**: pulling knowledge out of memory strengthens it far more than re-reading it (Roediger & Karpicke, 2006) — the reason every module in this course opens with a 60-second primer. **Spacing**: distributed practice beats massed practice. **Interleaving**: mixing problem types forces discrimination — which method applies here? — rather than autopilot repetition. **Elaboration**: explaining ideas in your own words and connecting them to what you already know — the reason every lecture in this course ends with a one-sentence explanation task.

Notice what these four have in common: each one is *inefficient* from the student's point of view. Each feels slower and more error-prone than the alternative. And each is exactly the kind of work a fluent text generator will happily do instead — producing the summary before the student has retrieved, the explanation before the student has elaborated, the worked solution before the student has discriminated. AI does not threaten learning by being wrong. It threatens learning by being *helpful at the wrong moment*.

> **PAUSE AND APPLY — find the effort**
>
> Take the assignment you held in mind in Module 1. Find the single component where students do the most effortful processing — the place where they retrieve, discriminate, explain, or revise rather than collect and format.
>
> Now ask: *if a student handed that component to AI, would anything visible change?* If the answer is no, you have just located your assignment's Safety Gap — and the likely site of your Cognitive Sanctuary.

Productive struggle is not needless friction

Here is the distinction the whole module turns on. **Productive struggle** is friction that builds the schema: working out an argument from primary sources before reading a critic; attempting the proof before seeing the worked example; drafting the interpretation before receiving feedback. **Needless friction** is friction that blocks access without building anything: an inscrutable assignment prompt; formatting labor that precedes any feedback on substance; a hidden prerequisite nobody teaches; a required tool that assumes a laptop, a subscription, and a quiet evening. Both are friction. Students experience both as difficulty. Only one is pedagogically valuable (Module 2 lectures).

The distinction matters because *friction is not the same as busywork, and rigor is not the same as difficulty*. A course can be brutally difficult and build nothing; a course can feel supportive and demand

retrieval at every turn. The diagnostic question for any assignment component is: **is this friction the thing students should learn to do, or is this friction in the way of what students should learn to do?** If the former, protect it — deliberately, visibly, and with a stated rationale. If the latter, remove it — and removing it is where AI earns an honest place in your course.

AI is excellent at removing needless friction: reformatting references, restructuring a document, translating instructions, explaining a dense passage at an accessible reading level, summarizing logistics. These are real wins, and refusing them on principle mostly punishes the students with the least slack. But AI is equally excellent at removing productive friction — generating the argument before the student has tried, elaborating before the student has elaborated. The tool does not know the difference. The syllabus has to.

> **COMMON TRAP — two mirror-image mistakes**
>
> **Trap 1 — "All difficulty is rigor."** Defending every hard part of an assignment as if hardness itself were the point. Formatting a bibliography in an archaic style before receiving any feedback on the argument is difficulty; it is not learning. Rigor theater drives students toward AI shortcuts and falls hardest on students with the least time.
>
> **Trap 2 — "All friction is cruelty."** Smoothing every difficulty out of the path — or letting AI smooth it — until the course is frictionless and nothing sticks. A frictionless course is not a kind course; it is a course that quietly exports the struggle to some later semester, employer, or licensing exam where the stakes are higher and the support is gone.

Naming the thinking students must still do

"Students must still think" is too vague to design around — the Module 1 lesson about "critical thinking" applies here too. The Friction Audit therefore asks you to name the essential cognitive operation from a fixed six-item menu: **retrieve** (produce from memory), **discriminate** (tell cases, sources, or methods apart), **explain** (account for why, in your own words), **apply** (use a concept on a new case), **critique** (evaluate against criteria), and **revise** (improve work in response to feedback). One operation, named precisely, is worth more than a paragraph of aspiration: it tells you what the Cognitive Sanctuary must protect, and later it tells Module 3 what the assessment must actually measure.

The audit also asks you to write the assignment's outcome as a *capability*, not a product. "Students submit a five-page essay" names a product. "Students discriminate between two scholarly framings of the same event and defend a choice" names a capability. Products can be generated; capabilities have to be demonstrated. The rest of the worksheet — and frankly, the rest of the course — runs on that substitution.

Cognitive Sanctuaries and Access Repairs

The Friction Audit ends in two design moves, and the pairing is deliberate: one protects struggle, one removes barriers, and the course refuses to let you do only one of them.

A **Cognitive Sanctuary** is a bounded course moment where students do essential thinking without AI substitution — not an AI-free course, but an AI-free *moment* with a stated purpose. Workable forms include a brief in-class or LMS-quiz explanation in the student's own words; an annotation pass over sources the student selected; a sixty-second recorded walk-through of a decision; a process note naming the alternative the student considered and rejected. Good sanctuaries are small (minutes, not hours), explained to students (the rationale is part of the pedagogy), and aimed precisely at the essential operation you named — not at everything.

An **Access Repair** is a barrier you remove because it is not pedagogically valuable: an inaccessible file format, an ambiguous prompt, a hidden prerequisite, excessive formatting labor, an assumed paid tool. The worksheet asks how AI may legitimately help with the repaired barrier — and what the alternative is for students who cannot or choose not to use AI. That last field is not boilerplate; it is the difference between an accommodation and a new requirement.

EQUITY AND ACCESS CHECK

Needless friction is not evenly distributed. The same barriers that are minor annoyances for well-resourced students — dense prompts, rigid formats, assumed devices, assumed quiet time — fall hardest on multilingual students, first-generation students, working adults, and students with disabilities. Removing needless friction is an equity move, not a softening move (cf. CAST's Universal Design for Learning guidelines).

The same logic runs in reverse for the sanctuary: if your AI-free moment assumes fast handwriting, same-day campus presence, or spoken English under time pressure, you have protected the struggle for some students and created needless friction for others. Design the sanctuary with your least-resourced student in the room.

Posture, briefly: matching AI's role to the learner's stage

Module 1 introduced Oracle, Adversary, and Tutor as stances toward AI. Module 2 sharpens them into design choices matched to learner stage. The **Oracle** (fast, fluent answers) is usable by experts who can evaluate output, and it is the default posture when no other is specified — which is why unspecified "AI allowed" usually means Oracle. The **Tutor** (asks questions, withholds answers, gives graduated hints) preserves productive struggle while reducing isolation, which suits novices — but only with explicit boundaries; a tutor that yields the full answer after one wrong attempt is an Oracle in costume (Promptbook; Module 2 reference card). The **Adversary** (AI produces output; the student critiques it) builds evaluation capacity, but only for students who already have enough schema to find the flaws. The most common design error in AI-era assignments is posture mismatch: putting a novice in front of an Oracle, or asking a novice to critique errors they cannot yet see. The Friction Audit's component table is where mismatches become visible — and the matrix below is the sorting tool.

	AI could remove this friction	AI cannot (yet) remove this friction
Friction builds the intended learning (productive struggle)	Highest risk — protect it. Design a Cognitive Sanctuary around the essential operation; say why in student-facing language.	Protect and monitor. Capability drift can move this cell left — recheck each term (vintage: June 2026).
Friction blocks access without building learning (needless friction)	Repair it. Let AI remove the barrier; document the change and provide a non-AI alternative path.	Redesign it. The barrier serves no one and AI will not fix it — simplify the prompt, format, or logistics yourself.

Table 1. The friction-sorting matrix behind the Friction Audit. Every assignment component lands in one cell; the cell tells you the design move. (Adapted from Module 2 lectures and worksheet.)

Connection to the Module Artifact

CONNECTION TO YOUR PORTFOLIO ARTIFACT — Friction Audit (with one Cognitive Sanctuary and one Access Repair)

The Friction Audit worksheet walks the path this reading just walked: name the assignment and its outcome-as-capability; mark each component P (productive struggle), N (needless friction), or B (both); name the essential cognitive operation from the six-item menu; mark where AI could remove it; then design one Cognitive Sanctuary and one Access Repair.

Use the same assignment you will carry into Modules 3 and 4 — the audit becomes the diagnostic evidence for your AIAS decision and your redesign. If you are working from a Synthetic Course Packet, the worksheet's worked example (Introductory Statistics) shows the expected grain size.

Before you submit, confirm:

- □ My outcome statement names a *capability* (verb a student performs), not a product (thing a student submits).
- □ Every major component is marked P, N, or B — and at least one honest N appears (an audit with no needless friction usually means I stopped looking).
- □ The essential cognitive operation comes from the menu: retrieve, discriminate, explain, apply, critique, or revise.
- □ I marked where a generic AI tool could remove that operation (capability claims dated June 2026).
- □ My Cognitive Sanctuary specifies what students do, what they submit, the time budget, and how it is evaluated — and it targets the essential operation, not everything.
- □ My Access Repair names the barrier, why it was not pedagogically valuable, how AI may legitimately help, and the alternative for students who cannot or choose not to use AI.

Key Terms

Term	What it means in this course
Safety Gap	The distance between what a student can produce with AI assistance and what they could independently do without it. A visibility problem before it is a policy problem.
Working memory	The small, limited mental workspace where new information is processed; the bottleneck that makes novices novices (Sweller, 1988).
Schema	Organized long-term knowledge that lets experts process efficiently. Built through retrieval, error, and elaboration — not by reading polished examples.
Desirable difficulty	A learning condition that feels harder in the moment and produces more durable learning — retrieval, spacing, interleaving, elaboration (Bjork & Bjork, 2011).
Retrieval practice	Strengthening memory by pulling knowledge out rather than re-reading it; the reason each module opens with a 60-second primer (Roediger & Karpicke, 2006).
Interleaving	Mixing problem or case types so students must discriminate which approach applies, rather than repeating one type on autopilot.
Elaboration	Explaining ideas in one's own words and connecting them to prior knowledge; the reason each lecture closes with a one-sentence explanation task.
Productive struggle	Friction that builds the schema — the difficulty that is the learning.
Needless friction	Friction that blocks access without building learning; removing it is an equity move, not a softening move.
Cognitive Sanctuary	A bounded, explained course moment where students do essential thinking without AI substitution.
Access Repair	A barrier deliberately removed because it is not pedagogically valuable — with a documented non-AI alternative.
Posture mismatch	Assigning an AI role that does not fit the learner's stage — most commonly, a novice in front of an Oracle.

Brief Applied Example

A statistics instructor sorts the friction. Professor Okafor teaches introductory statistics to sixty students, many of them first-generation and most of them anxious about the material. His weekly problem sets ask students to choose a test, run it in a spreadsheet, and write a short interpretation. On the audit, the outcome-as-capability comes out as: *given a research scenario, discriminate which statistical test applies and explain the choice* — discriminate and explain, two operations from the menu.

The component table is clarifying. Choosing the test: P, and a current AI tool produces a defensible choice with reasons in seconds (claim dated June 2026) — highest-risk cell. Running the computation: B, leaning N; the arithmetic was never the point. Formatting the output tables to match a house style: N, pure friction. Writing the interpretation: P. His Cognitive Sanctuary becomes a weekly two-question LMS quiz, closed-book and low-stakes, where students see a *new* scenario and write two sentences: which test, and why — five minutes of protected discrimination, graded credit/no-credit. His Access Repair removes the formatting requirement entirely and explicitly permits AI for cleaning up tables and rephrasing instructions, with a documented spreadsheet-template alternative for students who do not use AI tools. The problem sets stay; what they prove has been honestly relocated. When a colleague asks whether the quiz makes the course harder, Okafor has the module's answer ready: it makes the course *visible* — the struggle was always there; now it is in the right place.

End-of-Reading Reflection

1. Return to your 60-second primer response. Is the struggle you named productive struggle or needless friction — or did writing the distinction change your guess about *why* the struggle is happening?
2. For your own assignment (or your Synthetic Course Packet's), name the one component you would defend to a skeptical colleague as essential friction — and the one you would quietly delete tomorrow.
3. The reading claims AI threatens learning by being "helpful at the wrong moment." Where in your course is the wrong moment? What would make it visible to students as the wrong moment, rather than merely forbidden?

Sources and Further Reading

The module materials cited throughout this reading — the Module 2 lectures, the Module 2 Participant Workbook, the Oracle/Tutor/Adversary reference card, the Safety Gap diagram, and the Friction Audit samples library — are provided inside the course.

Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher et al. (Eds.), *Psychology and the real world*. Worth.

CAST. (2018). *Universal Design for Learning guidelines* (version 2.2). https://udlguidelines.cast.org

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio/Penguin. Chapter 7.

Promptbook: The power of educational research, generative AI, and openness. (n.d.). EdTech Books. https://edtechbooks.org/promptbook.

Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science, 17*(3).

Stanford Teaching Commons. (n.d.). *Artificial intelligence teaching guide*. Stanford University. https://teachingcommons.stanford.edu/teaching-guides/artificial-intelligence-teaching-guide

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science, 12*(2).

Suggested optional extension: the Promptbook's Socratic-tutor chapters preview the bounded Tutor prompts you will adapt in Module 4; Mollick's chapter 7 is the most readable single treatment of why "the homework apocalypse" is a design problem rather than a moral one.