

# Teaching with Artificial Intelligence

## MODULE 4

# Redesigning Assignments with TILT and AI-Aware Scaffolding

*Faculty Development Reading*

*Design for the learning you need to see, not for the tool you hope to outrun.*

— the course stance, in one sentence

June 2026

## Module Purpose

---

Module 4 is the build module. Your Module 3 Decision Memo named a level, an evidence gap, and a planned change; this module converts that memo into a full, student-facing **Redesigned Assignment**. The design commitment is stated up front, because it shapes every choice that follows: this course does not promise AI-proofing. It helps you build assignments where AI use is *intentionally restricted*, *intentionally integrated*, or *intentionally critiqued* — and where the evidence of learning is robust whichever case holds (Module 4 lectures).

The module also draws a line the AI-policy literature too often blurs: this is **redesign, not surveillance**. You will add transparency, one piece of process evidence, and explicit AI-use guidance. You will not add prompt-log dragnets, detection rituals, or paperwork that mistakes volume for visibility. The major decision you are preparing to make: *how do I rewrite this one assignment so students understand what they are learning, what AI use is acceptable, and how success will be evaluated — without exploding anyone's workload?* The deliverable is the rewritten assignment plus a 100-word design rationale.

## Learning Objectives

---

After completing this reading and the Module 4 activities, you will be able to:

1. Explain why AI-proofing is fragile design and define AI-resilience as the more durable goal, in terms a department colleague could repeat.
2. Rewrite one assignment in TILT structure — Purpose, Task, Criteria — with an explicit AI Use section naming permitted, prohibited, and (if any) required uses, plus a documentation expectation.
3. Add exactly one process-evidence element that reveals a learning-relevant decision, and justify in writing why it is worth the workload it adds.
4. Design one AI-integrated or AI-critical move (Tutor, Adversary, or Collaborator) with a posture-preserving boundary and an alternative path for students who cannot or choose not to use AI — if and only if integration serves the outcome.
5. Audit the redesign for student workload, faculty review time, accessibility, and clarity.

## Before You Read: A 60-Second Retrieval Primer

---

**YOUR 60-SECOND RESPONSE — write before you read**

**Name one thing you would try first if asked to make your assignment AI-resistant. Then ask yourself: would your move still work if a new model came out next month?** Sixty seconds, one or two sentences. Keep your answer nearby — the first section of this reading is addressed directly to it.

---

## Core Reading

---

### AI-proof is the wrong goal

Nearly every faculty member's first redesign instinct is some version of AI-proofing: find the task the machine cannot do, and assign that. The instinct is understandable — it promises that current course materials can survive untouched — and it fails in four predictable ways (Module 4 lectures). **Capability drift**: today's AI-proof task is next term's AI-easy task, and the arms race always favors the side that ships updates monthly. **False confidence**: "AI can't do this" usually means "I have not yet tried" — the jagged frontier cuts both ways. **Inequitable enforcement**: when proofing fails, instructors catch the students they already suspect, not the students who used AI. **Brittle task design**: AI-proofing tends to generate arbitrary constraints — obscure formats, in-jokes, gotcha requirements — that block access without building learning; needless friction, manufactured on purpose.

The durable goal is **AI-resilience**: an assignment designed so that students may or may not use AI, and the evidence of learning is robust to that assumption. Resilience does not mean everyone uses AI. It is fully consistent with AIAS Level 1 — *no AI here, for stated learning reasons* — and incompatible only with what the course calls *secret Level 1*: assignments whose rules say "no AI" while the design quietly assumes students could not use it anyway. Three strategies produce resilience, and you already hold two of them: make cognitive work visible in process (your Module 2 Cognitive Sanctuary); specify AI use clearly enough to be documented (your Module 3 AIAS contract); and choose an authentic, discipline-specific context that generic AI output cannot satisfy. The module's compass is the **redesign triangle** — learning goal at the apex, AIAS level on one base corner, evidence form on the other. When the three corners reinforce each other, the assignment stands; when one is chosen for comfort rather than alignment, it wobbles.

### TILT: the chassis of the rewrite

The rewrite itself uses the Transparency in Learning and Teaching framework (TILT; Winkelmes), whose research base is refreshingly modest in its demand: tell students the **Purpose** (why this assignment matters to their learning — capability, not compliance), the **Task** (what they will actually do, in steps a tired reader can follow), and the **Criteria** (what quality looks like, concretely enough to guide work before submission). Transparent design was an equity intervention before AI made it urgent: the implicit expectations it surfaces were always easiest for insiders to guess. Module 3's lecture showed the before-and-after: a case-analysis prompt that read "analyze the case using course concepts" becomes a purpose (practice making a recommendation under constraint), a task (a two-page memo to a named program director, two course concepts, one cited piece of case evidence), and criteria (explains the decision; uses two concepts accurately; cites case evidence; names one tradeoff). Nothing was made easier. Everything was made visible.

To the TILT triplet the course adds a fourth labeled section: the **AI Use box**. It answers, in student-facing language: what AI use is *permitted* here, what is *prohibited*, what (if anything) is *required*, what documentation is expected, what privacy caution applies (what should never be pasted into a commercial tool), and what the alternative path is if AI access is uneven. The box does for AI what

criteria do for quality: it replaces guessing with a contract. A student who reads “You may use AI to brainstorm planning questions; the recommendation and reasoning must be yours; document any planning use in one sentence” knows exactly where the line sits — and why, because the Purpose paragraph just told them what capability the assignment builds.

#### COMMON TRAP — surveillance dressed as evidence

The failure mode of process evidence is process *paperwork* — and its extreme form is surveillance. Requiring students to submit every prompt and every AI response generates volume no instructor will read, exposes private data, normalizes monitoring, and drives honest students underground while changing nothing for dishonest ones (Module 4 lectures).

The test for any process requirement: *does this reveal a learning-relevant decision that would otherwise be invisible?* A decision memo that names the choice and the alternative considered passes. A folder of screenshots fails. If you find yourself designing a compliance trail rather than a window into thinking, return to the Stress Test’s question 2 and ask again what you actually need to see. And keep the Student AI Use Note short — the course’s template caps it at two paragraphs, because long documentation requirements are where honesty goes to die.

### Making thinking visible without making work explode

The redesign adds **exactly one** new process-evidence element. This is the module’s *rule of one*, and it is load-bearing: it forces you to choose the single piece of evidence that reveals the most learning-relevant decision, and it protects both your students and you from the redesign becoming a second course. If your current assignment already has several weak process requirements, consolidate them into one strong one — subtraction is a legitimate redesign move.

The menu from Module 3 still applies: an annotated source trail; a short decision memo; a sixty-second recorded oral defense; an LMS quiz with short-answer follow-up; a prompt log with reflection (only where AI is actually integrated); a student-version-versus-AI-version comparison with critique. Each candidate then faces the **workload audit** — two sliders, both honest: student time per submission, and faculty review time per submission multiplied by enrollment. The course supplies a calculator worksheet, but the arithmetic is the easy part; the discipline is believing it. Sixty seconds of review across sixty students is an hour per assignment cycle. If either slider crosses what your actual semester can absorb, the evidence is wrong-sized — choose a lighter form, not a heroic intention. A rubric line completes the element: strong process evidence *names the choice and an alternative considered*; weak process evidence restates the assignment.

### When AI joins the assignment: posture, boundary, and prompt

If your Module 3 decision was to integrate AI — and only then — the integration must specify three things, or “AI permitted” silently defaults to Oracle. First, the **posture**. A *Tutor* asks questions, withholds answers until the student has shown attempts, and suits Levels 2–3: the submission is the dialogue plus a one-paragraph reflection on what shifted. An *Adversary* generates plausible-but-flawed output for the student to critique, suits Levels 3 and 5, and submits the critique, never the AI output. A *Collaborator*

works iteratively with the student — draft first, suggestions second, student judgment last — suits Levels 3–4 and advanced students, and submits the final work plus a short decisions note (Module 4 Posture Prompts card; Promptbook).

Second, the **posture-preserving boundary**: one explicit rule whose violation signals the posture has collapsed, written into the assignment. The course’s model is candid with students: *if the AI gives you the answer before you have shown two attempts, paste the moment into your submission and note what happened — that is documentation, not a penalty*. A Tutor that answers on the first wrong attempt is an Oracle in costume; an Adversary the student cannot push back against is Oracle output with a villain’s name. Third, the **prompt itself**, supplied to students rather than left to chance. The course scaffolds prompts with iCRAFT — Imagine the AI’s role; Clear instruction; Relevant context; Audience; Format; Task/purpose (Urand, 2024, in the Promptbook) — and the Module 4 reference card provides starter prompts for all three postures, each adaptable by swapping in your discipline’s context and each workable on a free tier. Every integrated design also writes the **alternative path**: what the student who cannot or will not use AI does instead, at equivalent learning value and workload. If no integration serves the outcome, the worksheet asks for one sentence naming the alternative cognitive work — and Level 1 with a stated rationale remains a complete, respectable answer.

#### PAUSE AND APPLY — draft the one element

Before opening the redesign template, draft your single process-evidence element in three sentences: (1) the form, from the menu; (2) the learning-relevant decision it reveals — finish the sentence “this shows me the student chose...”; (3) the two workload numbers: student minutes, and faculty review minutes times enrollment.

If sentence 2 stalls, the element is paperwork; pick again. If sentence 3 frightens you, lighten the form now — not in week six.

### Accessibility and clarity are design requirements, not finish work

A redesigned assignment that cannot be understood at midnight by a tired, anxious, or second-language reader has not been redesigned; it has been relocated. Write the Task in short sentences and numbered steps. Keep the AI Use box scannable — labeled lines, not a paragraph of hedges. State time expectations honestly. Ensure every required material is screen-reader accessible and every required tool has a free-tier path; if the assignment involves AI-generated media, caption it. The course’s accessibility floor is pragmatic: a B2-level reader on a phone should be able to determine what to do, what is allowed, and what good work looks like, on the first pass. (Module 5 will run your whole policy through this same test; building the habit at assignment scale makes the policy nearly write itself.)

#### EQUITY AND ACCESS CHECK

Run the redesign past four students before you call it done: the student without a paid tool (is every required use free-tier feasible — and stated so?); the multilingual student (does any evidence form punish accent or prose speed rather than thinking? is there an equivalent alternative?); the student with a disability accommodation (does the process evidence interact with extended time, screen

readers, or speech input?); and the working student (can the process element be completed asynchronously, in fragments, on a phone?).

And one more constituency: the colleague who will grade it. An adjunct teaching three sections cannot review sixty oral defenses a week. If the design only works for the instructor who built it, it fails the course's portability standard before Module 6 even names it.

## The 100-word rationale — and the sentence students will misread

Two small artifacts close the module. The **design rationale** (100 words, instructor-facing) names one change you made and why — the seed of the cover memo your Module 6 portfolio will carry. And the reflective exercise asks you to find the one sentence in your redesigned assignment that students most need to understand correctly, then write down the likely misreading — the specific way a tired or confused student would take it wrong — and either revise the sentence or defend it. It is the cheapest usability test in the course: you already know where your assignments get misread; this exercise just makes you say it before the semester does. The module's three-layer self-review (rubric, samples library, and a five-question AI critical friend that probes misunderstanding, paid-tool assumptions, ambiguity, workload, and second-language barriers — without proposing edits) closes the loop.

| Step                          | What you do                                                                                                                    | What it produces                             | Carried from            |
|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|-------------------------|
| <b>1. Carry forward</b>       | Copy the outcome, AIAS level, and Stress Test finding into the redesign template.                                              | The design brief at the top of the page.     | Modules 2–3             |
| <b>2. TILT rewrite</b>        | Write Purpose, Task, and Criteria in student-facing language.                                                                  | The transparent assignment core.             | Module 3 (TILT preview) |
| <b>3. AI Use box</b>          | Name permitted / prohibited / required uses, documentation, privacy caution, and the access alternative.                       | The visible AI contract.                     | Module 3 (AIAS level)   |
| <b>4. One process element</b> | Choose one evidence form; write its rubric line; run the workload audit.                                                       | The visibility plan — rule of one.           | Module 2 (Sanctuary)    |
| <b>5. Posture + boundary</b>  | If integrating AI: posture, iCRAFT prompt, boundary, alternative path. If not: one sentence on the alternative cognitive work. | The bounded AI move (or the stated Level 1). | Modules 2–3             |
| <b>6. Rationale + review</b>  | Write the 100-word rationale; self-review against rubric, samples, and critical-friend questions; revise once.                 | The finished redesign for the portfolio.     | Module 4                |

Table 1. The TILT-plus-AI redesign workflow, as sequenced in the Module 4 Assignment Redesign template. Six steps, one assignment, one new process element.

## Connection to the Module Artifact

### CONNECTION TO YOUR PORTFOLIO ARTIFACT — Redesigned Assignment + 100-word design rationale

The Assignment Redesign template walks the six steps in Table 1 with example completions throughout and a full worked example from the Composition 101 synthetic course. Everything in this reading maps to a template section — the triangle to Section 2’s carry-forward fields, TILT to Section 3, the AI Use box to Section 4, the rule of one to Section 5, postures to Section 6, the workload audit to Section 7.

Write the assignment in genuinely student-facing form: this exact text should be able to drop into your LMS unedited. It is the central artifact of your final portfolio and the object your Module 5 policy must agree with.

### Before you submit, confirm:

- ☐ Purpose names the capability the assignment builds — not “because it’s required” — and echoes my Discipline Statement.
- ☐ Task is numbered, concrete, and survivable by a tired B2-level reader on a phone.
- ☐ Criteria are observable enough to guide work *before* submission, not just to justify grades after.
- ☐ The AI Use box names permitted, prohibited, and required uses; a one-to-two-sentence documentation expectation; a privacy caution; and a free-tier or non-AI alternative.
- ☐ Exactly one new process-evidence element exists; it reveals a named decision; both workload numbers are written down and survivable.
- ☐ If AI is integrated: posture named, iCRAFT prompt supplied, posture-preserving boundary stated, alternative path written. If not: the alternative cognitive work is named in one sentence.
- ☐ The 100-word rationale names one change and why — and I found and fixed (or defended) the sentence students will misread.

## Key Terms

| Term           | What it means in this course                                                                                                                                      |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AI-proofing    | Designing assignments AI supposedly cannot complete. Fragile: defeated by capability drift, false confidence, uneven enforcement, and brittle constraints.        |
| AI-resilience  | Design in which students may or may not use AI and the evidence of learning is robust either way. Consistent with AIAS Level 1; incompatible with secret Level 1. |
| Secret Level 1 | An assignment whose rules say “no AI” while the design quietly assumes students could not use it anyway — prohibition without resilience.                         |

| Term                                          | What it means in this course                                                                                                                        |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Redesign triangle</b>                      | Learning goal, AIAS level, and evidence form — three corners that must reinforce each other for a redesign to stand.                                |
| <b>TILT</b>                                   | Transparency in Learning and Teaching (Winkelmes): state the Purpose, Task, and Criteria of every assignment in student-facing language.            |
| <b>AI Use box</b>                             | The fourth labeled assignment section: permitted, prohibited, and required AI uses, documentation, privacy caution, and access alternative.         |
| <b>Process evidence vs. process paperwork</b> | Evidence reveals a learning-relevant decision; paperwork generates volume. The test is the decision, not the artifact count.                        |
| <b>Rule of one</b>                            | Add at most one new process-evidence element per assignment; consolidate weak elements rather than stacking them.                                   |
| <b>Workload audit</b>                         | Two honest numbers: student time per submission, and faculty review time per submission multiplied by enrollment.                                   |
| <b>Posture-preserving boundary</b>            | The one explicit rule whose violation signals the AI posture has collapsed — e.g., a Tutor answering before two attempts.                           |
| <b>iCRAFT</b>                                 | The Promptbook's prompt scaffold (Urand, 2024): Imagine the role; Clear instruction; Relevant context; Audience; Format; Task/purpose.              |
| <b>Student AI Use Note</b>                    | A short (two-paragraph maximum) documentation template: what AI was used for, what it was not, one decision made differently, any erroneous output. |

## Brief Applied Example

*A biology instructor redesigns the lab report.* Professor Whitfield's introductory biology students write five lab reports a term, and her Module 3 Stress Test was unambiguous: a current model produces a clean report from a photographed data sheet in minutes (claim dated June 2026), and the product alone cannot show the capability her course exists to build, *explaining why the result supports one conclusion rather than its rival*. Her Decision Memo chose Revise, Level 2, with the discussion section flagged as the evidence that mattered.

The rewrite goes through the six steps. Purpose tells students the report is practice in reasoning from evidence to claim, the move biologists are actually paid for. Task rennumbers the sections and shrinks the methods transcription she never read anyway. The criteria name four observable qualities, including "the discussion considers and rules out one alternative explanation." The AI Use box permits AI for formatting tables, checking calculation steps, and rephrasing the introduction (documented in one sentence); prohibits AI drafting of the discussion section; warns students never to paste classmates' data or personal information into commercial tools; and supplies a formatting template alternative for students who skip AI entirely. Her one process element is a **decisions note**: three sentences at the end of the discussion naming the alternative explanation considered and the evidence that ruled it out —

thirty seconds to read, sixty students, one hour per cycle, survivable. No AI posture is integrated; her sentence on alternative cognitive work reads, “The discussion is where you reason without assistance, because that reasoning is what the course certifies.” Her misreading exercise catches exactly the sentence she feared: students will read “AI may check calculation steps” as “AI may do the calculations.” She revises it to name the boundary — *AI may verify steps you have already written* — and her 100-word rationale records why. Total new faculty workload: one rubric line and an hour of reading per cycle. Total new student clarity: the entire assignment.

## End-of-Reading Reflection

---

1. Return to your primer response — your first AI-resistance move. Was it AI-proofing or AI-resilience? Would it survive next month’s model release, and if not, what in this module replaces it?
2. For your own assignment or Synthetic Course Packet: which TILT element is currently weakest — Purpose, Task, or Criteria — and what is the one sentence you would add to fix it?
3. Where is the line in your redesign between evidence and surveillance? Write the one process requirement you decided *not* to add, and why — that subtraction is part of your design rationale.

## Sources and Further Reading

---

The module materials cited throughout this reading — the Module 4 lectures, the Assignment Redesign template, the Posture Prompts reference card, the Student AI Use Note template, the Workload Audit calculator, and the redesigned-assignment samples library — are provided inside the course.

Harvard University. (n.d.). *AI teaching resources*. <https://www.harvard.edu/ai/teaching-resources/>

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

*Promptbook: The power of educational research, generative AI, and openness*. (n.d.). EdTech Books.  
<https://edtechbooks.org/promptbook>. Includes the iCRAFT prompt-design framework (Urand, 2024).

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

Winkelmes, M.-A. (n.d.). *Transparency in Learning and Teaching (TILT Higher Ed)*. <https://tilthighered.com>

*Suggested optional extension:* the Promptbook’s cognitive-apprenticeship chapter (“From Model to Mentor”) deepens the Tutor posture’s design logic; the Module 4 Posture Prompts card is worth printing and keeping beside your LMS.