

Teaching with Artificial Intelligence

MODULE 6

Portfolio Integration, Maintenance, and the Six-Month Plan

Faculty Development Reading

A portfolio is finished when a colleague could teach from it tomorrow. A plan is realistic when it fits inside ninety minutes a semester.

— the course stance, in one sentence

June 2026

Module Purpose

Five modules have produced five working artifacts. Module 6 does two unglamorous, essential things with them: it assembles them into a **teaching-ready portfolio** a colleague could pick up and use, and it builds the **Six-Month Implementation Plan** that keeps the portfolio alive after the course ends. The module's load-bearing standard is deliberately humble — the **substitute-instructor test**: could a colleague hand this packet to a substitute on day one of next term, and would the substitute know what to do? If yes, the portfolio passes. If no, it does not matter how elegant the rationale is (Module 6 lectures).

The second purpose is protection against time. AI capability drift will, left alone, quietly convert your good design into stale policy — the example that ages out, the access assumption that flips, the “AI cannot do this” sentence that becomes false over winter break. The major decision you are preparing to make: *is this packet ready to teach from, and what small, scheduled practice will keep it true?* Note the word small. The course's answer budgets ninety minutes a semester — a maintenance habit, not a second job.

Learning Objectives

After completing this reading and the Module 6 activities, you will be able to:

1. Assemble a course redesign portfolio that passes the substitute-instructor test, with a 100-word cover memo and a four-item readiness checklist.
2. Distinguish which AI expectations require department- or program-level alignment from those that should remain course-level decisions.
3. Triage your portfolio's elements by capability-drift exposure — high-drift, moderate-drift, stable — and run a Stale Example audit.
4. Create a six-month revision plan naming three triggers, two evidence sources, one department-level question, and one steward or community touchpoint.
5. Reflect on how your Module 1 Skeptic's Charter has been honored, challenged, or refined across the course.

Before You Read: A 60-Second Retrieval Primer

YOUR 60-SECOND RESPONSE — write before you read

Picture a colleague on day one of next term, holding your portfolio, about to teach your course. What is the first sentence they would need to read to know what to do? Sixty seconds. That sentence — or your struggle to find it — is the cover memo this module will have you write.

Core Reading

From six artifacts to one teaching packet

The portfolio is a packet, not a portfolio in the credentialing sense — the course is explicit about what it is *not*: no reflective essay on your AI journey, no bibliography, no abstract. Those genres serve evaluators; this document serves you, in week one, at 8 a.m. Its architecture has two layers. The **student-facing layer** is what students will actually encounter: the Redesigned Assignment (Module 4), the Syllabus AI Policy (Module 5), and the First-Week Touchpoint. The **instructor-facing layer** is the reasoning that makes the student-facing layer defensible: the Discipline Statement and Skeptic’s Charter (Module 1), the Friction Audit (Module 2), the Stress Test and Decision Memo (Module 3), the Equity/Ethics Audit (Module 5), and the Six-Month Plan, which this module adds. Keeping the layers distinct is not filing tidiness; it is what lets you hand the right pages to the right reader — students get clarity, colleagues get reasons, and you get both.

Two short pieces bind the packet. The **cover memo** — 100 words, five fields — names the course, the redesigned assignment and its AIAS level, the most important equity safeguard, and the element most fragile to capability drift, then asks for your signature and date. It is the sentence from your retrieval primer, grown up. The **readiness checklist** asks four blunt yes/no questions: syllabus-ready? LMS-ready? grading-ready? student-ready? Any “no” is not a failure; it is this week’s to-do list, surfaced now instead of in week one.

The substitute-instructor test

The module’s self-review replaces the samples library of earlier modules — by Module 6, the only relevant comparison is between this packet and your own next term. You read the portfolio as if you were a colleague handed it cold, and answer four questions in writing: Could I read the syllabus AI policy and tell students on day one what is expected? Could I grade the redesigned assignment without further instructions? Could I run the first-week touchpoint without modification? And — the question that finds the gaps — *what single piece of context is missing that a substitute would need?* If question four surfaces something critical, it becomes a one-paragraph addendum to the cover memo. The test reads as logistics and works as epistemology: it is a falsifiable standard for whether your design thinking actually made it onto the page, or still lives in your head (Module 6 lectures; cf. the OpenAI education newsletter’s recurring emphasis on human review and documented guardrails in AI implementations).

PAUSE AND APPLY — run the test on one artifact now

Do not wait for the full packet. Take your strongest artifact — most likely the Module 4 assignment — and run question two against it: could a colleague grade this without you in the room? Read the criteria as a stranger.

Write down the first thing the stranger would have to ask you. That question marks the exact spot where your design is still tacit — and one sentence, added now, is cheaper than one email thread per student later.

Reopening the Skeptic's Charter

In Module 1 you signed a charter naming what AI must not damage in your teaching, what you would accept as evidence of damage, and what you would accept as evidence of non-damage — and you were asked not to edit it since. Module 6 reopens it, unedited, and asks for one paragraph beneath: what was **honored**, what was **challenged**, what you **still hold**. The discipline of the exercise is symmetry. If the course gave you reasons to relax a fear, the paragraph should say which evidence did it. If the course sharpened a fear, say that with equal candor — the charter was never a conversion device, and a participant who ends the course with a firmer, better-evidenced “no” to some AI uses has used the course exactly as designed. What the paragraph cannot honestly say is “nothing changed,” because at minimum you now hold named operations, levels, and evidence forms where there were instincts. Skepticism toward instruments is a different professional posture from skepticism toward mood.

Departmental alignment without forced uniformity

Your portfolio governs one course; your students live across many. A student carrying three courses with three contradictory AI regimes will resolve the contradiction in whatever way maximizes survival, not learning — and the inconsistency itself becomes an integrity hazard. The opposite failure is just as real: a blanket department rule that flattens disciplinary and developmental nuance, forcing the advanced workshop and the introductory survey into the same contract for tidiness's sake (Module 6 lectures).

The course's resolution: **align on vocabulary, decide at course level**. Three things are worth aligning across a program: the shared AIAS vocabulary, with examples, so the same words mean the same things in every syllabus; a shared minimum equity floor — no required paid tiers, no detection scores as standalone evidence; and a shared, restorative-first handling of suspected violations. Three things should stay with the course: the specific AIAS level per assignment, the specific posture (Tutor, Adversary, Collaborator, or none), and the specific process-evidence form. The distinction gives you language for the department meeting you may be dreading: a colleague can keep an all-Level-1 course and you a Level 3 project — students are served not by uniform rules but by *legible* ones. The module provides a five-scenario **department scenario poll** you can run asynchronously and anonymously; its job is not consensus but visibility — surfacing the variance that already exists so the alignment conversation starts from facts rather than vibes.

Capability drift: triage before maintenance

Capability drift was named in Module 1; Module 6 finally answers it. The answer begins with triage, because not everything in your portfolio ages at the same speed. **High-drift** elements are claims tied to current model behavior: specific tool names, free-versus-paid feature boundaries, detection-reliability claims, any sentence of the form “AI cannot do X.” Expect these to wobble within months (this reading's own capability claims: June 2026). **Moderate-drift** elements shift on institutional time: policy language tied to contracts and provisioning, course-level AIAS assignments as norms evolve. **Stable** elements are the frameworks themselves — TILT, the AIAS as vocabulary, post-plagiarism, the Safety Gap, the friction distinction. They are stable because they describe learning rather than software, which is also the quiet vindication of the course's design bet: spend the judgment on what lasts (Module 6 Capability Drift

Watchlist; cf. Mollick, 2024, ch. 9, on planning under the assumption that today's model is the worst you will ever use).

Triage feeds the **Stale Example audit**, the plan's most concrete chore: walk the syllabus and the assignment, find every concrete claim about what AI can or cannot do, check its vintage, and re-verify or rewrite anything older than six months. It is twenty minutes of work that prevents the specific embarrassment with the highest pedagogical cost — a confident classroom claim that a student's phone can falsify from the back row.

The Six-Month Plan: triggers, evidence, and a calendar block

The plan itself is one worksheet section, and its fields are designed to be filled in an evening: **three revision triggers** — events that prompt a look, drawn from a menu including a major model release, a change in institutional AI policy, student questions or violations revealing a policy gap, end of semester, and a department conversation; **two evidence sources** you will actually consult — student artifacts from the term, a re-run of one Stress Test question against your assignment, the course's annual Capability Drift Bulletin, departmental input; **one department-level question** you will bring to a colleague, and **one steward or community touchpoint**. The plan closes by scheduling a single **90-minute calendar block** before next term — the entire maintenance budget. The constraint is the design: if your plan demands more than ninety minutes a semester, the redesign was over-engineered, and the plan will be abandoned rather than executed. Maintenance you will actually perform beats vigilance you will only intend.

When	What you do	Trigger you are watching	Evidence you consult	Time
Month 0 (course end)	Assemble portfolio; sign cover memo; schedule the 90-minute block; set the Future-Self Letter reminder.	—	Readiness checklist; substitute-instructor test.	This module
Months 1–2 (term begins)	Teach with the packet. Run the first-week touchpoint; note student questions verbatim.	Student questions revealing a policy gap.	Touchpoint responses; help-seeking patterns.	In course
Month 3 (mid-term)	Reopen the Future-Self Letter; check one prediction against reality.	The evidence you said you would look for.	Student artifacts from the redesigned assignment.	~20 min
Months 4–5 (term ends)	Skim the term's artifacts against your Friction Audit; flag one thing that worked and one that did not.	End of semester (standing trigger).	Graded work; one Stress Test question re-run.	~30 min
Month 6 (before next term)	The 90-minute block: Stale Example audit; revise one policy or assignment paragraph; check the Capability Drift Bulletin; bring one question to the department.	Major model release; institutional policy change.	Drift Bulletin; department scenario poll results.	90 min

Table 1. The six-month maintenance timeline assembled from the Module 6 plan fields. The total scheduled cost after course end is roughly two and a half hours.

COMMON TRAP — the heroic maintenance plan

The most common Module 6 failure is writing a plan for the person you wish you were in January: monthly tool reviews, a reading group, a full assignment re-validation every break. That plan is abandoned by February, and its abandonment quietly discredits the whole portfolio — *see, none of this was realistic*.

The repair is subtraction. One calendar block. One Stale Example audit. One revised paragraph. One department question. If you finish early, wonderful — but design the plan for a tired person in a busy week, the same reader standard your policy already meets. The matching trap at the other extreme: treating the portfolio as a finished monument. It is versioned, not finished; the cover memo's date is there because a second date will join it.

Sustainable agency: what this course did and did not do

It is worth being precise about what the portfolio certifies, because the course refuses to oversell it. It did not make you an AI expert; expertise in a moving field has a shelf life measured in months. It did not settle whether AI is good for your discipline; your Skeptic's Charter response paragraph is allowed to end unconvinced. What it built is a *repeatable judgment practice*: name the capability, sort the friction, choose the contract, design the evidence, audit the equity, schedule the revision. That loop — not any single artifact in the packet — is the durable thing, because it is the part that survives the next model release. Faculty agency in the AI era will not come from outrunning the technology; it comes from owning a process that does not need to outrun it (Mollick, 2024, ch. 9). The **Future-Self Letter** — one page to yourself three months out, naming one thing you expect to work, one you expect to fail, one piece of evidence you will check, and one revision you are willing to make — is the loop's first scheduled turn.

EQUITY AND ACCESS CHECK

The substitute-instructor test and the 90-minute budget are equity commitments, not conveniences. A portfolio that only its author can teach from concentrates AI judgment in the tenured, the course-released, and the long-tenured — and leaves the adjunct hired in August to improvise a policy alone. Teaching readiness is the rubric's *first* criterion precisely to protect the colleague with the least slack.

The same lens applies to your department question: programs whose AI vocabulary lives in five veteran faculty members' heads have not aligned anything. A shared AIAS vocabulary, a shared equity floor, and a packet a newcomer can teach from are how AI judgment becomes program property instead of personal property.

Connection to the Module Artifact

CONNECTION TO YOUR PORTFOLIO ARTIFACT — Final Course Redesign Portfolio + Six-Month Implementation Plan

The Portfolio Template provides insertion slots in order: cover memo; table of contents; Discipline Statement (with a one-sentence revision prompt if your Module 1 thinking has shifted); Skeptic's Charter plus your new response paragraph; Friction Audit; Stress Test and Decision Memo; Redesign Assignment (confirm student-facing form); Syllabus AI Policy with Equity/Ethics Audit; First-Week Touchpoint; the Six-Month Plan fields; and the readiness checklist.

Assembly is mostly insertion — the thinking is done. Reserve your attention for the four pieces this module adds: the cover memo, the Charter response paragraph, the Six-Month Plan, and the substitute-instructor test with its addendum if question four finds a gap. The example completed portfolio in the Module 6 folder shows the target grain size.

Before you submit, confirm:

- ☐ All artifact slots are filled, and the student-facing pieces (assignment, policy, touchpoint) are in genuinely student-facing form.
- ☐ My cover memo is ~100 words and names the course, the assignment and its AIAS level, the key equity safeguard, and the most drift-fragile element — signed and dated.
- ☐ My Skeptic's Charter response paragraph names what was honored, what was challenged, and what I still hold — with the evidence for each.
- ☐ The Six-Month Plan names three triggers, two evidence sources, one department-level question, and one steward/community touchpoint — and the 90-minute block is on my actual calendar.
- ☐ I ran the substitute-instructor test in writing; question four's missing context (if any) became a cover-memo addendum.
- ☐ The readiness checklist is honestly answered — and any "no" has a date attached.

Key Terms

Term	What it means in this course
Teaching-ready portfolio	The assembled packet of all six artifacts, organized for use in week one — not a reflective or credentialing document.
Student-facing vs. instructor-facing layers	What students encounter (assignment, policy, touchpoint) versus the reasoning that defends it (statement, audits, memos, plan).
Substitute-instructor test	The portfolio's load-bearing standard: could a colleague teach from this packet on day one without you in the room?

Term	What it means in this course
Cover memo	One hundred words naming the course, the redesigned assignment and AIAS level, the key equity safeguard, and the most drift-fragile element.
Readiness checklist	Four yes/no questions — syllabus-ready, LMS-ready, grading-ready, student-ready — that convert gaps into a dated to-do list.
Capability drift triage	Sorting portfolio elements into high-drift (model-behavior claims), moderate-drift (institutional language), and stable (the frameworks).
Stale Example audit	Finding every “AI can/cannot do X” claim, checking its vintage, and re-verifying or rewriting anything older than six months.
Revision trigger	A named event that prompts a portfolio review: model release, policy change, student questions, end of term, department conversation.
Six-Month Implementation Plan	Three triggers, two evidence sources, one department question, one community touchpoint, and one scheduled 90-minute block.
Shared vocabulary, course-level decisions	The department-alignment principle: align AIAS language, the equity floor, and restorative handling; leave levels, postures, and evidence forms to each course.
Future-Self Letter	A one-page letter to yourself three months out: one expected success, one expected failure, one piece of evidence, one revision you are willing to make.
Capability Drift Bulletin	The course’s annual steward update on what changed in AI capability and which stale stamps need refreshing.

Brief Applied Example

A studio art instructor assembles the packet. Professor Calder teaches introductory studio art, eighteen students, and spent the course working on his observational drawing portfolio assignment — Level 1 for the drawings themselves (the looking is the learning), Level 5 for a new critique unit in which students compare AI-generated images against their own studies and articulate what the model gets wrong about light, weight, and decision. His portfolio assembles cleanly until the substitute-instructor test, where question four catches him: nothing in the packet tells a substitute how to run a critique of AI images without the conversation collapsing into better/worse. The missing context becomes a cover-memo addendum: three sentences naming the critique’s actual axis — *what did the human study record that the generated image cannot account for?*

His drift triage is instructive because studio art sits close to the fastest-moving capability frontier. High-drift: every sentence about what image generators render poorly — hands, typography, consistent light — stamped June 2026 and first in line for the Stale Example audit. Moderate-drift: the campus license question (which tools students may use without paying) that he routes to his department question. Stable: the Level 1 rationale, because no model release changes why drawing from observation builds seeing. His Six-Month Plan watches one trigger with special attention — a major image-model release —

and his Future-Self Letter predicts one failure honestly: he expects two students to submit generated images as studies, and commits to treating the first instance as a restorative conversation about what the assignment certifies, exactly as his Module 5 policy promises. The 90-minute block lands on the Tuesday before spring term. His Charter response paragraph ends: *“Still skeptical. Better instrumented.”* The steward, reading it, will recognize a course completed as designed.

End-of-Reading Reflection

1. Return to your primer sentence — the first thing a substitute would need to read. After this module, is it still the right first sentence? Draft the cover memo’s opening line now, while the gap between your head and the page is fresh.
2. Reread your Module 1 Skeptic’s Charter before opening the portfolio template. Which of your own predictions are you now most curious to test next term — and which piece of evidence from Modules 2–5 moved you most?
3. For your course or Synthetic Course Packet: which single element of your portfolio will go stale first, and what is the calendar date on which you will check it?

Sources and Further Reading

The module materials cited throughout this reading — the Module 6 lectures, the Portfolio Template and example completed portfolio, the Capability Drift Watchlist, the Department Scenario Poll, and the Future-Self Letter worksheet — are provided inside the course.

Mollick, E. (2024). *Co-intelligence: Living and working with AI*. Portfolio/Penguin. Chapter 9.

OpenAI. (n.d.). *ChatGPT for education* newsletter archive. <https://edunewsletter.openai.com>.

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

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge University Press.
(Background for the course’s community-of-practice and critical-friends framing.)

Suggested optional extension: Mollick’s closing scenarios pair well with the drift-triage exercise — read them with your watchlist open; the annual Capability Drift Bulletin is the course’s standing answer to “how will I know what changed?”