Applying Anchored Adaptation

How to build AI literacy at your institution when no fixed plan can stay current.

For the convener, the unit that runs the work · Leo S. Lo, 2026

The idea in one line. Anchored. Then: do, capture, iterate. You commit to a framework and to a decision each cycle; you refuse to freeze the implementation, because in a fast-moving field a frozen plan is obsolete before it scales.

This playbook is the companion to the definition (v1.0), maintained separately as a living document. The definition states the principle; this shows how to apply it. It is written for the convener; the academic leaders who sponsor the work have a companion of their own, the provost playbook.

THE SIX STEPS

The definition states the method as one commitment and a three-beat loop: anchored, then do, capture, iterate. Step 1 sets the anchor; step 3 is where you do, step 4 where you capture, step 5 where you iterate. Steps 2 and 6 are the institutional and human conditions that make the loop possible.

1 **Set the anchor.** Adopt or adapt one framework that defines AI literacy in durable, tool-independent terms. Test a candidate first: a sound anchor is tool-independent, teachable, assessable, durable across at least one technology cycle, and broad enough to span the literacy domain. Once set, hold it. Review the anchor itself only on a long cadence, annual or longer, and change it only on explicit evidence. If the anchor is ever revised, every standing decision bound to it comes back for review, so revision marks a boundary in the record rather than a loss of it. The anchor is what keeps adaptation from becoming drift.

2 **Find the convener.** In practice, change in a decentralized institution moves bottom-up, so implementation needs a home: name one unit with cross-school reach, convening legitimacy, and no territorial stake, often the library. Give it a light coordinating role: a convener, not a mandate. This step is institutional practice rather than part of the method's definition, but without it the loop has no host.

3 **Run small, real experiments.** Start where there is energy, not authority. Take genuine teaching, research, or operational problems and solve them with AI, coached from question to evidence. In the first cycles the coaching comes from the convening unit (a librarian or teaching-center consultant is enough), and coaching capacity is itself something the cycles build. Mix tempos: semester-long course work alongside short sprints.

4 **Capture every cycle.** Require each engagement to leave a durable artifact: an assessment, a documented case, a template, a student portfolio piece. For each consequential decision you scale, fund, or stop, keep a consistent record: at minimum, the choice made, the evidence behind it, and a date to revisit it. That minimum record is enough to start; note in it the anchor version in force, so a later anchor revision can find the decisions it touches. The evidence stays local, formative, traceable, and bound to a decision; it documents how the institution chose, not what an intervention caused. Capture is the institution's memory, and what separates learning from churn.

5 **Feed decisions, then re-decide.** Name the owner: a sponsor or small decision group (the academic leaders from whom the convener draws legitimacy) receives the evidence on the cadence and owns what to scale, fund, or stop. The convener coordinates; the sponsors decide. Decide on current evidence, set a revisit date, and revisit. A reversible decision does not need causal proof;

plausibility, local fit, and a revisit date carry it, and that lower evidential bar is exactly what reversibility buys. Reversible decisions on a rhythm, not one fixed plan.

6 **Build the people, not just the process.** Because the work is coached, it also builds clarity, capacity, and culture, the human conditions that decide whether any of it holds. Treat adoption as change management, one engagement at a time.

WHERE FIXED STRATEGY WINS

Anchored Adaptation suits questions that are genuinely unknown and reversible; the definition calls these limits its boundary conditions. Some questions instead demand a settled answer, and the method does not apply to them.

Use a fixed answer for	Why iteration fails here
Legal and compliance	FERPA, data privacy, and records obligations require a settled position, not a cycle.
Procurement and contracts	Vendor terms commit the institution on a fixed timeline with fixed language.
Accessibility	Access is owed in full from the start, not earned over passes.
Safety-critical or irreversible	Where a wrong cycle causes harm that cannot be undone, learning by doing costs too much.
Hard external deadlines	Accreditation and regulatory dates leave no room for an open revisit cadence.

> **Rule of thumb: adapt where the answer is unknown and the decision is reversible. Fix where the answer is externally determined or the decision cannot be undone.**

The method also tells you when to retire it. A question has settled when successive cycles stop producing evidence that changes the decision. When that happens, write the fixed answer, stop cycling on that question, and move the loop to the questions that remain open. The method applies while implementation outruns governance; its own cadence shows you when that condition ends and fixed strategy resumes.

WHAT YOU NEED TO START

- One framework as the anchor.
- One convening unit with a coordinating role.
- A sponsor or small decision group that owns scale, fund, and stop.
- A few willing faculty and a small amount of flexible funding; its chief use is buying faculty time.
- A simple way to capture artifacts and a cadence to review them.

This does not require large staffing or deep internal AI expertise to begin. It requires a framework, a convener, and the discipline to capture and review.

WHAT TO AVOID

- Writing a comprehensive AI strategy or policy first. It dates before it clears governance.
- Launching big instead of learning small.
- Running pilots that produce no reusable artifact.
- Treating it as a tools problem rather than a judgment-and-culture problem.
- Deciding once. The cadence of re-deciding is the method.

WHAT TO WATCH

Three signals, reviewed on cadence: every engagement leaves a durable artifact; faculty return for a second cycle; and captured evidence begins entering real decisions about what to scale, fund, or stop. Measures of literacy gains and program outcomes come later, once enough cycles accumulate to assess them. The method claims only what its evidence can show. Watch also for the settling signal above: cycles that no longer change decisions mean that question is ready for a fixed answer.

STARTING CADENCES

Defaults, not rules: adjust to your calendar, and let the cycles revise them. A typical cycle runs four to twelve weeks, with course-length work alongside; a sensible first cohort is three to six experiments, small enough to coach well; the decision review meets each semester, or quarterly where the rhythm allows; and the anchor is reviewed annually or longer. Expect the first Watch signals (artifacts and returning faculty) within two or three cycles; if they have not appeared after a year, treat that as evidence and revise the approach.