

The Campus Education Harness

Executive Summary

The central challenge of artificial intelligence in relation to learning is “cognitive offloading,” where the learner delegates the task of learning to AI. However, this is not the only option. While one use of AI can diminish the “productive struggle” that is essential for learning, it is also possible for AI to serve as a partner in learning and transformation. The Campus Education Harness is constructed to orchestrate that partnership so that student learning is amplified rather than diminished.

This system stands between the learner and artificial intelligence, and we call it an “Adaptive Learning Agency Orchestration System.” It ensures that the learning process is optimized for each individual learner. Combined with AI, this system creates an “Educational Intelligence Platform” which maximizes student learning and transformation. Research is clarifying that access to artificial intelligence does not, in itself, improve learning. What matters is how AI is designed to partner in the learning process.

What is an Adaptive Learning Agency Orchestration System, as we have constructed it?

It is **adaptive**. That is to say, it adjusts its approach in real time to optimize the learning for the individual student.

It is about **learning**. It is about learning on every level, not simply knowledge but skills and disposition as well. That is to say, it aims at transformation. It is optimized for the student, not for the professor or some abstract sense of the subject. If the student is not appropriately changed, then teaching has not succeeded.

It is about **agency**. It is not about relinquishing the learning to a machine. It is about amplifying learning by facilitating a situation where the learner takes agency for their own learning. It does not give answers. It pushes the learner to exercise their intellectual muscles, only telling enough to lead the learner forward.

It is an **orchestrating** mechanism. Like the conductor of an orchestra, the platform cues the different elements of learning in the optimal order for the individual student. It integrates the components of learning into a whole that is suited for each learner as an individual in a recipe cooked for one.

Lastly, it is a **system**. It is a “harness.” A harness is a mediator between artificial intelligence and an individual learner. It takes the seemingly endless resources potentially provided by AI and “harnesses” it in a way that 1) reflects the specifics of an institution, that is, the context of the learning and 2) stays within the boundaries of the organization’s identity.

In what follows, we walk through how and why the harness does this. Key is that the harness knows both your institution and the individual student. This context enables the system to know where the student is on their journey both in relation to individual subjects and their broader path through your institution. The system knows what has worked with each individual student in the past, as well as what approaches to learning have not. From moment to moment, the harness knows what a student needs next. It knows the right tool for that step, and it knows the difference between a nudge and an answer.

At the same time, each institution has its own brand and its own boundaries. The harness filters the input of large language models so that the student engages what is most appropriate for each institution. The institution can trust that they, not AI or Campus, is determining the parameters of student learning and transformation.

We close with what this looks like for a single student, what it means for institutional control, and what it means for your institution going forward.

II. A Fork in the Road

Students are using AI. They are using it to do homework, to write papers, and to prepare for exams. They are using it with or without the permission of the instructor or institution. The question is, how are they using it? Are they using it in a way that “offloads” the productive struggle that is essential to learning, or are they using it to amplify and accelerate their learning in the most productive ways?

Research is beginning to suggest that the two paths lead to very different destinations. In one study of Chinese high school students, the use of an AI assistant increased homework scores by 18% and reduced completion time by 30%, while decreasing exam scores by 20% within six months. [1] This decline was concentrated among students who used AI to finish their work quickly rather than to engage with it. In short, the assistant had not improved learning. It had simply done the work for them. The resulting gap in learning then showed up in an exam room where the assistant was not present to help them.

This is the phenomenon of cognitive offloading. Rather than struggle through the thinking itself, the student delegates it to artificial intelligence. It is easy to miss while it is happening. The grades are good. Engagement seems high. The loss of learning is unseen until it is measured – or it remains unseen because it is never measured at all.

There is another way. It is harder to build, but it is the one that can actually use AI to amplify learning while giving faculty visibility into how their students are actually doing. It is a system that is constantly observing where a student is in the process of learning. It knows the knowledge, skills, and dispositions a student has mastered, and which ones are still developing. Such a system knows when and where to “nudge” without handing a student the answers.

We can build such a system so that it operates within the boundaries of what the institution affirms and what the professor teaches. The learner engages the material in a way appropriate

to the course and institution, not from a generic approach or from the standpoint of a conflicting version.

This is the fork in the road that every institution now faces. Students are already using AI. The question is whether their use quietly offloads the thinking that learning requires or whether we orchestrate that learning in a way that amplifies it. We can build systems that coordinate this partnership and thereby ensure that students remain engaged in the productive struggle that leads to real learning.

The Campus Education Harness exists because we believe the orchestrated path is the only one worth building toward. Reaching these goals requires more than good intentions. It requires the kind of system this paper describes.

[1] Strömberg, Lei & Wu, “The Generative AI Learning Penalty.”

III. The Seven Layers of the Harness

In software engineering, a harness is a mediator between artificial intelligence and an institution or individual learner. It is the structure that surrounds a system and directs it toward specific tasks. In educational technology, this mediation takes on a new purpose. It focuses AI's vast capacity and puts it to work within the parameters of an institution and its faculty. The seven layers that follow are what makes that possible.

Context

Context is what makes the system's intelligence usable rather than generic. Out of the vast array of information that AI can draw upon, context is what focuses it. Context knows the content that a specific course teaches and the learning outcomes of the course. It knows the mission of the institution, its foundational principles and perspectives. It also knows the specific approach to a topic that a particular professor takes.

Context keeps the responses of AI tethered to a specific institution and course, rather than floating free. While memory tracks what a particular student has done and where that learner is going, context grounds both the system and student in the identity of the institution and the course. This context stretches from the institution's overall mission down to the outcomes of a single assignment, and it knows what mastery of those outcomes actually looks like.

Orchestration

If Context is what the system knows, Orchestration is what the system does with it. This is the conductor. This is the layer that decides, moment by moment, 1) how to guide a student and 2) how much help to give them, and it constantly adjusts both as the student works.

Guidance is about direction. Does the system pose the next question and set the next task, or does it step back and let the student decide what to work on? Help is about assistance. When the student is stuck, does the system offer just a nudge, or does it hand over the answer?

A student new to a topic usually needs more guidance. They need clear direction on what to work on next. At the same time, they still need to do the thinking themselves, so AI should not give too much help. A student who has shown real command of the material needs less guidance, so they learn to direct their own learning. Yet they still benefit from help now and then, on the moments that are genuinely hard.

Proper orchestration knows how to strike that balance. It gives enough guidance to keep a student moving, and withholds enough help to keep them struggling productively.

Every response a student receives has passed through this decision first. And Orchestration does not simply trust that its own instructions were followed. It checks, after the fact, on every response, what kind of help was actually given, not just what was intended. A system that means to offer a small hint but drifts into handing over the answer records that error so the record stays honest. This is what lets the harness keep faith with its own purpose: not what the AI meant to do, but what it actually did.

Pedagogy

Orchestration decides what to do in the moment. Pedagogy decides what counts as learning in the first place. As it turns out, that is a harder question than it sounds.

Not every correct answer means a student has learned something. If a student only gets to an answer with significant help, that success belongs in large part to the AI. Perhaps a full explanation was given, a worked example, or a solution laid out for them.

The harness calls this “brittle mastery.” It is when a student looks capable but is actually standing on assistance that will not be there later in an exam room or a job. Pedagogy is what keeps the system honest about the difference. A student has to show real, independent command of a skill more than once before it counts as mastered. This is not because the harness does not trust a single good answer, but because mastery is not ensured on one try.

Pedagogy also has to figure out why a student is struggling, because the right response depends entirely on the cause. A student who has never seen the material needs different help than a student who has a specific wrong idea in their head. Both of these are different from a student who understands the material but is afraid to attempt it or a student who has simply stopped trying.

Treating all struggle the same helps some students but does nothing for the rest. It may just blindly give more hints or more explanations without true guidance. Pedagogy's job is to tell each obstacle apart and to call in the right response to each one.

Tools

A tutor who only talks is limited. A good tutor knows when to point a student to a textbook, when to hand them a calculator, and when to simply answer from what they already know. The Tools layer is what gives the harness that same range.

Left alone, an AI system will often just generate an answer from what it already contains. It does so even when a better answer exists somewhere else, perhaps a search engine, a calculator, or in the specific readings and materials a course has actually assigned. The harness knows what tools to reach for and when.

A question about the reading gets answered from the reading, not from the AI's general sense of the subject. A math problem gets checked with an actual calculation, not just a plausible-looking one. A live, current question gets a live, current answer.

This matters because the alternative is exactly the kind of generic response the harness is built to avoid. This is a generic AI that answers everything the same way, out of the same well of general knowledge. Just as Context keeps a student anchored to their own course and institution, Tools makes sure a student's answers and when they need help come from the right source rather than being manufactured from thin air.

Memory

Context knows the course. Memory knows the student.

Over time, the harness builds a picture of each learner. That picture is not just what they have gotten right or wrong. It is the shape of how they learn. It remembers what a student has actually mastered, as distinct from what they have merely gotten past. It remembers which kinds of help have worked for this student before, and which haven't. After all, a worked example might click for one student and do nothing for another. And it remembers how independently this student tends to work: whether they reach for help quickly, or push through struggle on their own before asking.

This is what allows the harness to meet a returning student as someone it already knows, rather than as a stranger each time they log in. A student does not have to re-prove basic mastery every session, and the system does not have to guess from scratch what kind of help will actually land.

That memory is what makes real adaptation possible. It is not a system reacting to the last five minutes but one responding to a student's whole history in the course—and beyond.

Memory also knows what belongs to a single moment and what does not. What a student is doing right now matters for the very next response. How much help have they needed on today's topic? Have they made a real attempt yet? However, it does not define who they are as a learner. That distinction keeps the system fair. A hard day doesn't get baked in as a permanent judgment, and a strong one does not erase the work still ahead.

Permissions

Every other layer in the harness depends on a simple premise. The institution, not the AI and not Campus, decides what happens inside its own walls. Permissions is where that premise becomes real.

An institution has a mission. A department has a set of outcomes it is accountable for. A professor has a specific way of teaching a specific course. Permissions is what makes sure the harness actually operates inside those lines. It ensures that that a student engages the material as this institution and this professor actually teach it, not a generic version of the subject, and not one that conflicts with what the institution affirms.

This is not a filter bolted on as an afterthought. It is upstream of everything else the harness does. Before Orchestration decides how to guide a student, before Tools reaches for a source, before Pedagogy judges what counts as mastery, Permissions has already set the boundaries those decisions operate within. The institution sets those boundaries once in a foundational layer, deliberately, rather than hoping every individual response happens to land inside them.

And because those boundaries are set by the institution itself, they are the institution's to change. What one college affirms, another may not. The harness does not carry a single, fixed set of values into every school it serves. It carries the values of the school it is serving.

Evaluation

A student's grade tells you how they did on one assignment. It does not tell you whether they are actually becoming a more capable, independent learner. And it doesn't tell an institution whether the tool it adopted is doing what it promised. Evaluation is the layer built to answer both of those harder questions.

For a student, evaluation means seeing whether the AI's role in their learning is shrinking over time, the way a good tutor's engagement should. As a student demonstrates real, independent command of a skill, the harness deliberately steps back. It offers less guidance on that question and less help because the goal was never a student who performs well with the system. It was a student who no longer needs it to achieve a certain competency. An institution can see this directly. It's not only whether scores are going up, but whether independent thinking is increasing as well.

For an institution, evaluation works at a different scale. When many students share the same incorrect idea about a topic, that is rarely a coincidence. It usually means something in the course itself needs to change, not the student.

The harness surfaces these patterns back to faculty. Which misconceptions are showing up across a class? Which explanations are actually landing, and where might the course material itself have a gap? A course a professor has taught for three years, with this kind of feedback, becomes measurably better at diagnosing its own students than it was in year one.

This is the harness closing its own loop. It does not simply claim to work. It shows an institution, in plain and specific terms, whether it is working, both for each individual student and for every student after them.

IV. Learner Illustration

It is 11:40 on a Tuesday night. Katie is in Module 6 of the Astronomy course, working through the reading on stellar death (embedded on the screen) before making her first post in the online discussion forum. The Campfire prompt box is sitting right there, waiting for any questions she might have.

She has one. “I don’t understand why stars eventually die.”

The harness already knows several things about this moment.

First, it has **Context**. This is Module 6, and the reading has laid out two distinct paths a dying star can take, depending on its mass. On one path, it ends as a white dwarf. On the other, it will supernova and end up either as a neutron star or black hole.

The discussion prompt due this week doesn’t just ask students to name the stages. It asks them to explain why a star’s fate depends on its mass. Later, in the same prompt, it asks them to engage the dominant theory holding that life on Earth only became possible after earlier generations of stars had already lived and died.

That second question isn’t neutral, and the harness doesn’t treat it as if it were. Because of **Permissions**, the system knows that this institution teaches the sequence of stellar death and rebirth as more than physics. It assumes that the universe is ordered with intention, the product of a Creator rather than chance.

When the harness responds to Katie tonight, it will not strip that framing out in favor of a purely naturalistic account, and it will not treat intentional design as an add-on to be balanced against a secular alternative. It teaches the science through the conviction the institution actually holds.

But the harness also knows the difference between what the institution has settled and what remains genuinely open — including questions scientists themselves are still working out. The exact mechanics of core collapse in the final seconds before a supernova, for instance, are still an active area of research; astronomers don’t fully agree on every detail of what tips a dying star’s core past the point of no return. If Katie had asked about that instead, the harness wouldn’t manufacture a false certainty the field itself doesn’t have. It would be honest about what is settled, and what is still being discovered — which is itself a stance the institution holds deliberately, not a gap in the harness’s knowledge.

The harness knows these sorts of specifics about Katie because of **Memory**. The harness knows these sorts of specifics about Katie because of Memory. The system also knows that she has done well before when given a question to work through rather than an explanation.

In this case, the system knows that Katie has a tendency to confuse a star's mass with its brightness. Accordingly, she doesn't understand why a star gets brighter even as its mass is decreasing. The harness knows that she showed this confusion on a quiz she took two weeks ago.

So **Orchestration** decides. Give guidance, not help. Instead of an explanation, Katie gets a question back.

"What do you think determines whether a dying star becomes a white dwarf or a supernova?"

It does not give the answer. It gives a nudge toward the exact distinction the professor already knew she would need.

Because of **Pedagogy**, that nudge is aimed precisely at Katie's likely confusion, rather than a generic "think about it more." And because of **Tools**, the question uses the same word the textbook does—mass, instead of a synonym that might not connect.

Katie tries again. She gets partway there. "Mass matters," she remembers, but she isn't sure why.

One more nudge, slightly more specific, still short of the answer. On her third attempt, Katie writes it herself.

"A star's mass determines its fate, because only a massive enough core can collapse hard enough to trigger a supernova."

It is 11:53 PM. She has an answer she can actually defend in the discussion. It is an answer that she arrived at herself in partnership with AI. She arrived at it through productive struggle, not through brittle mastery.

Evaluation records that this was real, independent work. It updates what it knows about Katie: the mass-versus-brightness confusion from two weeks ago is resolving. And when several other students hit the same wall in the days that follow, that pattern will not just disappear into the gradebook. It will reach the professor as confirmation of exactly what was predicted when the course was built. Next term, that declaration gets sharper still.

Katie never sees any of this. She experiences a slightly harder Tuesday night, and an answer that is actually hers.

V. Institutional Boundaries

Every institution that adopts the harness is trusting it with something significant. It is trusting it with the real conversations of its students, at all hours, on subjects that matter to the institution.

This trust only makes sense if the institution has laid the foundations for what happens inside the harness, not Campus and not the AI.

This is how the system works. The harness operates within the parameters that each institution has configured: its own mission, its own faculty declarations, its own sense of what a given course is for. The parameters of one institution have no impact on the parameters of another.

Two institutions with significantly different missions or ideologies, both operating with a Campus harness, are not running the same harness tuned differently. They are running with different harnesses, each shaped by the institution it serves.

Faculty hold real authority inside this framework, not symbolic authority. A professor who has taught a course for years knows things about it that go beyond surface level data. The professor knows where students get stuck, what a “wrong turn” usually sounds like, which questions are worth answering directly and which are worth turning back into a question.

The harness is built to take that knowledge directly as a standing instruction from the faculty. Eventually, the harness will gain those insights itself. What a professor declares, the harness checks for, and it checks for them every time, for every student, not just the ones a professor happens to notice is struggling.

And because these boundaries genuinely belong to the institution, they are also the institution's to revisit. A course gets redesigned. A policy changes. A professor decides a rule needs to be stricter, or looser, or different entirely. None of that requires touching a single line of code, and none of it requires Campus's involvement. The institution is not asking Campus for permission to run its own course. It is telling its own harness how to run it.

This is also where the institution retains oversight, not just influence. What the harness has decided, and why, is not hidden inside the model. It is recorded. It can be surfaced. And it is visible to the people responsible for the course.

If a pattern emerges across many students, faculty see it. If a student's engagement with the harness stalls out entirely, the system can flag that student for a human conversation rather than quietly continuing on its own. The harness is built to hand things back to people, not to keep them for itself.

None of this is possible with a generic AI tool bolted onto a course after the fact. It is only possible because the boundaries are built in from the start, as a condition of how the system works. These are not policies layered on top of it. They are not a promise the institution has to take on faith. They are integrated from the ground up.

VI. What This Means for Your Institution

None of this requires an institution to start over. A president or provost does not need to rebuild a curriculum, retrain a faculty, or adopt anyone else's values to put the harness to work. The

harness is built to operate on foundations that likely already exist. This could be a course a professor has taught for years under a mission the institution has long articulated with a set of academic policies already on the books.

What changes is not what the institution stands for. It is seeing these foundations integrated into the university's instruction on a level it has never imaged.

For a president, this is a question of differentiation and outcomes. Better than ever before, your institution can show, concretely, that its students are becoming more capable and more independent. You will have more detailed data than ever to show to families, accreditors, and to a marketplace increasingly skeptical of what a degree actually guarantees.

For a provost, the question of faculty trust is answered with a system that hands authority to instructors rather than working around them. For a CIO, the question of governance is answered. They have a system whose boundaries are configured, inspectable, and owned by the institution, not a black box making decisions no one can see.

Every institution is already living inside the fork this paper opened with. Students are already using AI, tonight, in a residence hall or a kitchen table, on an assignment due at midnight. The only real question is whether that use is quietly costing them the thinking they came to college to develop. Might rather it be maximized to partner with them on their behalf, deliberately, inside boundaries their own institution has set.

The Campus Education Harness exists to make the second outcome possible—not someday, but in the very next course an institution chooses to redesign. If that first course is one you already have in mind, that's where this conversation should start.

VII. Appendices