Mark Hopkins' Log

AI, Higher Education, and the Extraordinary Potential of an Ordinary Technology

By Matthew Pittinsky, Ph.D. · **DRAFT** *v24* · *June 2026*

In December 1871, at a Williams College alumni dinner at Delmonico's in New York, future U.S. President James Garfield rose to defend Mark Hopkins – the college's outgoing president – and gave American higher education an image it has invoked, arguably reified, ever since:

"The ideal college is Mark Hopkins on one end of a log and a student on the other."[1]

The quote has become something of a Rorschach test. To many, it represents the Socratic ideal of higher education – the charismatic instructor, the engaged learner, the intimacy of the instructional relationship, the truth of education as a social process. To critics it reads a bit more as static, elitist nostalgia, an unrealistic and unrepresentative picture of a higher education system that has always been more diverse in type and mission and practice than just the elite liberal arts model the quote represents.

For me, it's all about the log.

Garfield's toast was a rearguard action. At the dinner the alumni were openly weighing whether to push Hopkins out after thirty-five years — faulting his loose administration of the college and judging him too old, and too set in his ways, to keep pace with a fast-changing world; he stepped down within four months. Hopkins came to use "the log" to argue

1. In researching the quote, I learned it's possible Garfield said "log cabin" and not "log," which would have been a lot less poetic.

that the essence of a college was great teaching, not the accumulating apparatus of the modern university. To his dismay, as the historian Frederick Rudolph records, intellect "in the form of standards and libraries and scientific equipment" was gaining ground at Williams during his last years as president, and at his successor's inauguration he protested "the accumulation of materials and books" and made a last plea for the ideal of the log. Hopkins' successors built the libraries and laboratories he resisted.

More than 150 years later, in the digital age, higher education finds itself engaged in a similar debate. Now, most of us take for granted the physical infrastructure and tools that support teaching and learning but question the digital apparatus. (I write this essay a few weeks after the Learning Management System [LMS] Canvas went down during finals period, painfully illustrating just how central this digital infrastructure has become.)

While discomfort, if not resistance, has long existed with the centrality of digital platforms in how faculty "deliver" their courses, the volume is now dialed to 11. AI technologies are transforming what digital platforms can do for instructors and learners (and their relationship, including their power relationship). The debate is not just about the instructional value of these technologies, but their compatibility with the very values of higher education we hold dear.

Where I Sit

In 1997, I co-founded one of the first Course Management Systems (CMS), Blackboard, named the company ironically after the medium we meant to transform ("log" would have been an awkward brand), helped lead it as CEO and Executive Chairman until 2006, and now 20 years later

am excited to return and help lead it again beginning this November. I also have an academic background. When I left Blackboard, it was to complete my doctoral studies in the sociology of education, which I did, and then joined the sociology faculty at Arizona State University as a tenure-track assistant professor.

I believe deeply in the unique social organizational form of the university for creating and disseminating knowledge. I believe strongly in the value of a liberal arts education, especially in this technical age. I consider the relationships I developed with my students and the fostering of their "sociological imagination" – about the role of education in society, and how schools operate as organizations – to be my most cherished experience as a faculty member.

While it is fun to see Will Hunting (Matt Damon) take down a smug Ivy undergraduate, he is wrong when he says, "You dropped 150 grand on a fucking education you could have got for $1.50 in late charges at the public library." Most people need more than that, and higher education is indeed more than that.

While I have every commercial reason to tell you the LMS has developed into what Blackboard's early mission statement foretold – "a powerful environment for teaching and learning, revolutionizing education more profoundly than any innovation since Gutenberg created the printing press," it has not. Nor has it broken the "age-old barriers of access, quality, equity, and cost," another promise of our early mission, although you can make an argument in favor of access given the growth in online enrollments.

Somewhere along the way, despite the technology platform shifts of cloud and mobile, as well as decades of research and practice, the course

management system, while rebranded a learning management system, remains largely the same.

The Argument

I believe AI is a collection of ordinary technologies with core strengths — generation, assessment, and orchestration — that bring extraordinary capabilities to the future of the learning management system. Collectively, they enable differentiation and personalization, the holy grails educational technology has chased for decades. AI technologies are not going to magically transform education, nor are they the threat to the core of the instructional relationship that Hopkins saw in the apparatus of his day.

As I argued in The LMS at 30: Part Two, for thirty years, the digital platform at the center of higher education has been a course management system masquerading as a learning management system. AI technologies – their compelling capabilities – will force the transition from the former to the latter. The LMS can finally deliver on what I call the three Ls: organize itself around the Learner, actively support the activities of Learning, and hold a verifiable record of what was Learned.

Two consequences flow from this realization of a true LMS. First, we will see a massive shift in where pedagogical value lives supporting the instructor, from the publishing layer into the delivery system.

Second, the very AI technologies blamed for breaking learning and the assessment of that learning are uniquely positioned to save it. This is because the anxiety we feel right now is our sudden reckoning with what Stanford sociologist and historian David Labaree, building on the work of

Randall Collins, calls Credentialism. We've misdiagnosed the symptom as the disease.

AI will impact education at four levels – curriculum (what is taught), pedagogy (how it is taught), the delivery model (how a program is designed, produced, and priced), and the delivery system (the digital infrastructure through which all of it is organized and experienced). In this essay I focus on the delivery system as it reflects the shape and needs of the other three, as well as paces their realization.

So, how do we get from course management to learning management?

An Ordinary Technology

It begins with recognizing that AI is an ordinary technology that can do extraordinary things; emphasis on ordinary technology. At the 2026 ASU+GSV Summit I attended a panel called Every Student Should Have an AI Learning Companion. Early on, Isabelle Hau of the Stanford Accelerator for Learning pushed back on the title itself. Don't call it a "companion," she said. A companion is a human social concept. The panel never did define what it meant by a learning companion, but Hau's instinct was right, and it points to why the language we use for AI matters.

Mechanically, a modern AI system is a statistical model of patterns. A neural network is a very large set of numerical weights, tuned by exposure to examples, that maps an input to a probability distribution over possible outputs. There is no ghost in the machine. As Arvind Narayanan and Sayash Kapoor have argued, AI is best understood as a normal technology. By "normal technology" they mean a general-purpose tool that society

adopts gradually, over decades, as it did electricity and the Internet — not a separate, superintelligent species.

I describe its magic as pattern recognition at scale. Because the world we experience – whether words, images, sound, behavior, or physical matter – is fundamentally composed of "atomic" units assembled into patterns, a system capable of mapping, predicting, and generating patterns at a massive scale can do extraordinary things. A computer, after all, is just ones and zeros – electrical signals moving through etched sand – and look at everything it enables!

Ordinary in kind is not ordinary in degree. A normal technology can still exceed us — a car outruns any sprinter, a calculator any human arithmetician — and on many tasks that reduce to pattern recognition, AI already matches or beats us. That is the very reason the borrowed vocabulary of the human mind misleads: we do not say a car "runs," and we manage it as a machine rather than negotiate with it as a being. The more capable the technology, the more the words we choose for it matter.

Language Matters

We invite trouble when we describe AI technologies as something they are not. For the closing chapter of The Wired Tower, a 2002 book I edited on the Internet's impact on higher education, I invited the critic Neil Postman to contribute, drawn by the argument he had long made about technology. In Amusing Ourselves to Death, he had already distinguished a technology — the machine itself — from the medium, the social thing a culture makes of it. The technology "television" is the same box everywhere; television the medium is shaped by whether you have one channel or five

hundred, whether it auto-plays the next episode, and whether it runs commercials. The vocabulary we choose is how we build the medium.

One of the most consequential things a technology does is change our language. When we call a pattern-recognition model a “companion,” or say a system “thinks,” “knows,” or “reasons,” we are not naming a capability. We are lending the machine the words we reserve for human minds — and “to think” starts to mean what a model does, while “to learn” shrinks to what a model can produce. We gain a convenient vocabulary and lose the older meaning before we notice it is gone.

Students and faculty sense this. When the medium becomes a presence rather than a tool, in Postman’s term it turns mythic — accepted as natural, no longer interrogated — and it becomes a social actor in a way a tool never is: a presence can take your job; a presence changes what learning is. We made this mistake before, when “virtual classroom” obscured the difference between typing into a text box and live dialogue with a professor.

Strip the anthropomorphic language off, describe the capabilities as software, put the software in the hands of faculty and learners, and much of the apprehension turns to interest. Part of my confidence is the compelling nature of what AI technologies can do. Indeed, in the College Board’s February 2026 survey of more than three thousand faculty, sixty percent of those who used AI in their role as faculty agreed it could “help students understand complex concepts through personalized explanations.” That is about 1.7 times the share with *no* direct experience who agreed, and based on capabilities that are rapidly improving and largely bolted on.

An academic planner that tracks deadlines and models GPA scenarios is a genuinely useful piece of software. A study aid that generates flashcards

and level-appropriate explanations from your actual course readings is something students have always needed. A system that can draw on a professor's syllabus well enough to answer a question at 2 a.m. is support that scales. Taken together they are what I often call the Hogwarts Book of Magical Study Aids, risk of litigation by J.K. Rowling notwithstanding. Described as software, these are tools faculty can direct, students can use, and institutions can govern.

The Digital Log Emerges

The closest analog to AI's impact on education is the Internet's rise in the mid-1990s, which gave birth to the LMS. Last year marked the LMS's thirtieth anniversary, prompting me to write The LMS at 30: Part One (and then Part Two). Here is a short version of the history for context.

Before the Internet, the digital learning environment was the PC, and then the networked PC. PLATO, out of the University of Illinois in the 1960s, imagined networked instruction long before the Web. In 1985, a decade before the LMS, Steve Jobs founded NeXT to serve higher education. As described in Geoffrey Cain's book on NeXT, Jobs took the stage at EDUCOM to describe a "scholar machine," with students dropping into simulated worlds to learn everything from physics to DNA, to the court of King Louis. That said, it wasn't until the Internet that a true digital learning infrastructure took form at most colleges and universities around the world.

This is what made the potential of Blackboard so exciting. Even a 19th-century one-room schoolhouse had basic instruments for engaging in the learning process: the blackboard was the operating environment, chalk the development tool, the pointer a navigation device. What was the

Internet-based equivalent of this infrastructure? As I wrote in our 1996 concept paper, an instructional management system is, "in essence, an operating environment in which teaching and learning can occur." It was the digital log with all the affordances a digital, network environment offers.

Educom's Instructional Management Systems (IMS) project, also launched in 1996 with commercial partners and major research universities, was working toward the same idea. According to an FAQ at the time, IMS was "an industry specification project aimed at creating an open architecture that vendors and I/T departments can base the development of their software [on] to ensure compatibility and simpler integration." In other words: middleware for discipline-specific, pedagogically rich tools, unified around the learner and the instructor.

Five years later, in October 2001, Blackboard launched the Building Blocks program with an open SDK and a first cohort of partners ranging from Carnegie Mellon's Virtual Chemistry Lab to WebEQ for complex mathematical and scientific equations. IMS eventually became 1EdTech, which continues to champion LTI and much of the interoperability stack on which the modern EdTech ecosystem sits.

As an operating environment, leveraging open standards, I expected each digitally supported course to reflect its discipline, its instructor's pedagogy, and its institution's approach to learning. That is not how the market voted. Early applications like WebCT and Blackboard earned adoption by providing turn-key tools and features that helped manage course activities and enable Web publishing rather than rethinking how a digital environment could orchestrate learning.

By 2002, Gartner observed that "although billed as 'enterprise' products, today's mainstream course management systems are built primarily

around the course… their main function is managing course information, not supplying teaching and learning tools." Nevertheless, "Course Management" was rebranded "Learning Management," aligning higher education with the term used for comparable software in corporate learning. The acronym changed, but not the paradigm.

What would earn the L in LMS? The three Ls I noted earlier must be realized: the product must be organized around the Learner, used for Learning, and serve as the system of record for what was Learned. Today's LMS does the opposite on all three.

1. It is built around the course, not the learner – the home screen is a series of course tiles, not the student's academic planner.
2. Learning happens in the links out of it; the term I hear most often is that the LMS is a "container."
3. And what is left at the end of the semester is an enrollment and a letter grade – the organs of a record of learning are present, but the circulatory system that would carry evidence from an assignment up to a program outcome for accreditation and out into a credential for learner mobility is missing.

Each shift matters on its own. The category change comes from joining all three.

Sustaining and Disruptive

Against this backdrop of a static, course-centric LMS category, AI technologies represent both sustaining and disruptive opportunities. They are sustaining in the way they make the existing scope of the platform – the document publishing, the workflows, the click-after-click-after-click of course management – 10x better. Asking for a test bank that assesses the

philosophical tradition of Pragmatism drawn from your lecture notes is a far better experience than today's workflow. None of that changes the paradigm; it improves the current one. Where AI technologies are disruptive is in their ability to change the value calculation of the LMS itself, opening a new kind of value the platform never delivered: personalized learning services built on generation, assessment, and orchestration, centered on the learner, the learning, and the learned.

My rough estimate is that today's LMS is about seventy percent course publishing and management, twenty percent communication and collaboration, and ten percent instructional orchestration. In the AI era it is easy to imagine the weights flipping, with orchestration – the personalized learning services learners do not have today – becoming the primary value. That is the disruption: not a better container, but a different vector of value.

There is a paradox here. The very thing that made the LMS succeed — a turn-key course-management value proposition, simple enough that any instructor could stand up a course website without help — is what kept it from becoming a learning platform. Buyers kept paying for the container, requesting more items on the course management checklist, so the container is what developed (and overdeveloped), while the harder platform-and-orchestration value proposition went unrealized.

There is a name for this. Education historians David Tyack and William Tobin coined the "grammar of schooling" — grade levels, class periods, Carnegie units, teacher-centered instruction, structures so deeply embedded that reform efforts tend to accommodate them rather than change them. The LMS inherited that grammar and digitized it. We organized the software around the course because the course is the university grammar's basic unit; building it around the learner would have

meant changing the grammar, and the grammar usually wins. It is a big reason why an LMS from 1999 and one from today still look and feel the same. The LMS reflects the same as course management of old, because the course looks the same, which I guess is self-reinforcing, if not tautological.

We have an opportunity to reset. Incumbents – Canvas, D2L, Moodle, and yes, Blackboard – could botch this opportunity, bolting a chat window and workflow agent onto a course container and calling it transformation. There are customer-centric reasons they might look for the keys where the light is brightest and bank sustaining gains, which are real and immediate, without doing the more difficult, disruptive work of rebuilding the platform around the learner. The agentic LMS is not inevitable.

What the Platform Unlocks

The founding figure of the LMS category is, in my judgment, Carol Twigg, who ran Educom's National Learning Infrastructure Initiative – the home of the IMS project – and then founded the National Center for Academic Transformation. The central claim of her work is that you do not get the benefit of a new technology by simply digitizing instruction. You must redesign the course around what the technology enables – rethinking how we use lecture hours, TA and faculty time, and the ratios of who does what. NCAT's redesigned courses, on average, lowered cost by about a third while improving learning outcomes in 25 of 30 redesigns – cost and outcomes improving together, where they usually trade off. These were the large, introductory courses that enroll a large percentage of students.

Echoing Robert Dreeben and Rebecca Barr's work in K-12, Carol's work demonstrated that how we allocate a few core resources – resources

institutions actually control – has profound implications for teaching and learning: time, content, space, and people. The standard model fixes them in a handful of permutations – courses meet once or twice a week (plus lab), classmates are a function of prerequisites and degree path, class time is for introducing materials, etc. Most “innovation” merely digitized that arrangement instead of rethinking it.

I never asked her, but I’ve always assumed Carol’s leadership with IMS was based on her recognition that you need a digital learning platform to organize and enable the discipline and pedagogical tools at the center of redesigns. (And you need it to be standards-based.) A capable learning platform makes adjusting these resources more practical. Pace can be based on mastery. Time can be asynchronous as well as synchronous. Content can be differentiated. The division of labor between instructor, TA, learner, and peers can be facilitated differently. NCAT’s own redesigns showed it: Virginia Tech’s Math Emporium replaced lecture sections with an open, software-driven learning center where students worked at their own pace, with on-demand help from a tiered staff of faculty, graduate assistants, and peer tutors — time, space, and people reconfigured at once. Other redesigns kept the lecture but moved practice and feedback online, or replaced part of class time with it.

To be clear, the platform does not redesign the course; it makes redesign practical and affordable. Twigg’s results took deliberate institutional work, not software alone. What is new is not the idea – it is, for the first time, a delivery system capable of putting this kind of redesign within reach of an ordinary department, not just a well-funded pilot.

From Publisher to Platform

Writing about the potential of a true learning management system reminds me of one of my most memorable moments from the early days of Blackboard. It was 1999, and across the table sat Will Ethridge, then president of Pearson's Higher Education and Professional Publishing Group, with Jim Bankey and Pat Leonard. They wanted to license Blackboard to wrap course management around offerings like MyMathLab. Pearson was worried – reasonably – about being disintermediated. About the LMS becoming the aggregator of faculty and institutions, with the publisher reduced to what today might be called one tile in someone else's app store. (That we called third-party publisher offerings course cartridges did not help.) We agreed on a private-label deal in 2000, and Pearson invested in Blackboard as part of it, partly to be inside the tent of a potential disruptor.

I share this story because one consequence of the digital log I am describing is that the pedagogical value supporting faculty will move from the publishing layer to the delivery system. By pedagogical value I do not mean the instructor's own — that stays with the human — but the value that supports the instructor: the textbook's scaffolding, the worked example, the adaptive engine in courseware like MyLab, Connect, or ALEKS, which for thirty years has lived with the publisher. Publishers like Pearson and McGraw Hill have developed sophisticated offerings backed by robust performance data.

Two factors will shift that value to the learning infrastructure. First, many of the resources faculty once assigned at high cost can be produced on demand by the platform, drawing on the corpus of the course as raw material. Second, and more important, the algorithms that select what a learner needs at a given moment must operate at the platform level, not

the publisher level, to benefit from a student's full learning history rather than one chapter of one course. A typical undergraduate uses one publisher's eText in one course, another's homework system in a second, an open reader and the professor's own materials in a third; no single publisher covers a whole section, let alone the multi-year arc of a degree. The pedagogical layer is a many-to-many problem — many publishers, many tools, many courses, many learners — and a publisher-bound intelligence layer only builds silos.

A publisher-bound layer fragments the learner data model, so no one sees the full picture: faculty and institutions lose transparency over what is recommended to a student and on what basis, and personalization gets harder, not easier, because each engine sees only its own course. The capabilities that matter — a planner that holds the whole program, a study aid drawing on every course's corpus, support that knows where a student was strong two semesters ago — are not publisher-specific applications. Each requires an intelligence layer that crosses courses and operates on the full content a student or institution has licensed. Against today's scope of interoperability — gradebook passback, identity, rostering — what is being asked of the platform is an order of magnitude wider, and an order of magnitude more valuable.

I foresee two objections to this prediction of value shift.

First, an ordinary technology's impact on a domain depends on how digital that domain already is. AI is not a standalone invention; it is the latest layer on the PC, the Internet, and the cloud, and it transforms a field fastest where the work is already represented as data. Software engineering, fully digital, is changing almost overnight; medicine, rich in records and images, is moving fast but paced by the slower realities of clinical trials and regulation.

Higher education is unevenly digital, and that unevenness paces where AI lands. Because the LMS is still mainly course management, the sustaining gains — a better gradebook, faster content workflows — arrive first, while the deeper pedagogical change arrives more slowly. What the LMS as course management has digitized is largely summative assessment, not the formative work where learning happens. The richest learning data today sit not in the platform but with the publishers, whose courseware has long captured how students practice, struggle, and master a subject — which is why this shift will be gradual rather than sudden.

The shift is nevertheless inevitable. The unified model of the learner that real personalization requires is one only the platform can assemble, for the reasons above, and as the LMS takes on that orchestration layer the value will migrate to the delivery system.

Second, there has long been a belief that the LMS should remain trusted infrastructure, closer to a power grid than a pedagogical actor, with pedagogy left to discipline-specific tools built by expert authors. The LMS should be "pedagogically agnostic." I used to say that myself and obviously it needs to support a professor who simply teaches great literature by reading great literature, as well as a large biology course with many activities and forms of assessment. But when software can produce guided walkthroughs, graduated practice, and targeted feedback from any course corpus on demand, the question is no longer whether the platform will embody pedagogy. It will. The question is whether the institution gets to govern what the platform does, on what data, and under whose policies. Keeping the LMS "dumb" will not preserve neutrality; it will cede the pedagogical layer to whatever AI sits closest to the learner. The instinct to keep the platform predictable and well-bounded is the right one. The mistake is in thinking it requires keeping it inert.

The Great Unlock

When I taught SOC 421: Education and Society at Arizona State University, the assessments leaned heavily on writing – short papers, a longer team paper, in-class essay exams, regular discussion-board reflections. In a class of fifty or sixty students, the person who had to read all of it was me. Every faculty member who has graded a stack of fifty papers at midnight knows where this is heading.

In my syllabus I declared that the meta learning objectives of the course were to help students write well, speak well, think analytically, and be comfortable with numbers. Despite declaring these generalizable skills as important outcomes, looking back, the focus of my comments, and of the grade, was the substance: the argument, the evidence, the mastery of concepts. I did not engage with their writing in any depth. Reading my syllabus 16 years later, I am struck by its inconsistency. I directed students seeking help with drafts and revisions to ASU's Writing Center. The meta-objective got outsourced, not because I didn't care, but because I was one human, and the rubric in my head was richer than the time I had.

There is an old Silicon Valley adage that the difference between a product that gets adopted and one that fails is the difference between a vitamin and a painkiller. Assessment is the painkiller application of AI technologies – and not because AI created the pain. The pain was already there. The work of giving frequent low-stakes checks, of writing rich rubric-aligned feedback at the scale of a sixty or six hundred-student course, has always been difficult, and implemented less often than aspired.

For most of the last forty years, the critique of assessment in higher education has been the same: too much summative, too little formative. Paul Black and Dylan Wiliam's work pointed at feedback that helps a

learner close the gap between current performance and the standard she is reaching for. That pedagogy has been impractical at the scale of a real classroom, precisely because it required authoring and configuration work no faculty member had time for. Generative models, supervised by faculty and structured around their rubrics, may be the first software in the history of the LMS that can deliver this kind of assessment at scale. The pedagogy is not new. The capacity to run it is new.

The pedagogy experts have recommended for forty years turns out to be the pedagogy most robust to AI misuse. A take-home essay submitted at midnight is easy to fake. A semester-long sequence of small, formative checks – drafts with feedback, problem sets that build on yesterday's mistake, oral defenses, discussion that demands the student show her work – is much harder to fake, because the student has to keep showing up and growing in front of her instructor. How's that for a twist? The thing AI is blamed for breaking – student assessment – may be the thing AI is uniquely positioned to enable.

Closing the Gap

A student cheating with AI is behaving rationally; the system trained that rationality long before ChatGPT. For thirty years David Labaree has argued that American education serves three competing purposes – democratic equality, social efficiency, and social mobility – and that the third has steadily crowded out the others. Social mobility treats education as a private good: a zero-sum competition for the credentials that secure advantage. When the credential is the prize, the learning becomes incidental. If you are there for the grade and not the knowledge, Labaree observes, your incentive is to get it "the cheapest and easiest way I possibly can" – the effort a credential demands is just a price, and no

rational buyer pays more than they must. AI simply dropped the price to zero: the “magic sauce,” in his phrase, that lets a student “go through school without becoming educated.” AI did not create the gap between what a credential certifies and what a learner knows. It exposed it.

Seen this way, the two anxieties that dominate the AI conversation turn out to be one. The first is cheating: the student obtains the credential without the learning. The second is its mirror image: AI can lift a student who once earned C’s to A’s – but if the credential’s real job is to sort, and sorting requires that some fail, then everyone earning A’s is not a feature but a bug. Both fears are symptoms of the same condition: a system that has come to produce and protect the sort rather than the learning. We have misdiagnosed the symptom – AI – as the disease.

If that is the disease, the cure is not only to police the gap but to close it – to make the credential attest to learning again. And the same assessment machinery that defends against AI-enabled cheating is exactly what produces the evidence to do it.

Institutions need defensible answers to the questions their accreditors and boards routinely ask: what do your graduates learn and know, how well do they know it, what can they actually do? Run the assessment machinery well and it produces, as a byproduct of ordinary coursework, course-by-course evidence of what students actually learned, mapped to the objectives a program claims to deliver — not a parallel reporting burden on faculty, but broccoli in a smoothie: the vegetables go down with the strawberries, no separate plate required.

When I draft an assignment, the platform reads the syllabus, proposes which objectives the assignment tests, suggests rubric criteria, and I confirm or adjust. When a student submits, the rubric scores each criterion and the objective-level signal accumulates underneath the letter

grade. The student gets a paper back; the program gets evidence; the institution gets a record; nobody filled out a separate form.

We pitched a version of this at Blackboard twenty years ago — a system to "capture and report learning outcomes data while the learning is actually occurring," driving evidence up from course to program to institution. I spent years on the credential side of it at Parchment, where a transcript is supposed to attest to what a learner knows and too often attests only to a grade above C. That is the third L, Learned, made real. The same practices that defend against AI-enabled cheating may finally produce a credential that represents the learner.

There is a flip side. If the Internet's defining image was the 1993 New Yorker cartoon — "On the Internet, nobody knows you're a dog" — AI's might be this: "With AI technology, everyone is Captain Picard." (You need to picture it as a cartoon?) You direct, the technology executes. Picard's effectiveness comes from decades of judgment, experience, and knowledge. The Enterprise's computer is available to everyone. What you bring to it is what matters. For decades the honest answer to "why do I need to learn this?" was unconvincing, because the credential mattered more than the content — the triumph of exchange value over use value. AI changes the equation: substance becomes the source of advantage, not the credential divorced from substance.

None of this is automatic, and the same architecture that could close the gap could just as easily widen it. The dystopia closest to the instructional core is a system that does the assessment rather than supporting it: a faculty member divorced from the substance of student work, a student wondering why she is paying tuition to be evaluated by a model that has no stake in whether she learns. And a credential is only as trustworthy as our confidence that the work behind it is the enrolled student's own, not

software acting on her behalf — a problem policy and the platform will have to solve. The line between AI technologies an instructor directs in the context of software workflows and AI that replaces the instructor can't be blurry. We must draw it with purpose, in policy and in product.

Risks and Unintended Consequences

Which brings us to risks and unintended consequences. The 2026 Walton Family Foundation/GSV/Gallup survey of Gen Z found that excitement about AI fell fourteen points in a year, anger rose nine, and eighty percent think it likely that "using AI tools will make it more difficult for them to learn in the future." Faculty feel it too: in the AAC&U/Elon University survey — based on a non-scientific sample of faculty — ninety-five percent said generative AI will increase student overreliance and ninety percent said it will diminish critical thinking. This is a striking reversal. During the rise of the Internet, young people were the enthusiasts — Gallup and Pew found they overwhelmingly saw computer skills as essential to their future and expressed little fear of the risks. Now AI draws boos at commencement. The booing has many sources — job-market fear, an unlevel playing field, a distrust of the tech executives on the stage — and only some of them are about how we talk about the technology, as I suggested earlier.

Paraphrasing John F. Kennedy: there are real risks and costs to a program of *action* — but weighed against the forgone capabilities of AI, they are far less than the long-range risks and costs of comfortable *inaction*. That does not make the risks any less real. Two of Postman's premises were that all technological change is a trade-off, and that it always produces winners and losers. We owe a technology's effects as much scrutiny as we give its possibilities. The question is not whether AI in the LMS carries risks. Our

job is to name them and build for them — and different institutions and faculty will, and should, weigh them differently.

I was one of twenty-seven participants in a recent horizon scan of AI and higher education — researchers, university leaders, and people from industry — and the range of concerns and timelines was striking. Cognitive offloading is my top concern: if students hand the most difficult cognitive work to a model, the model does the work and the learning may not. The tendency to offload mental effort onto an external tool is well documented and precedes AI. The worry is that delegating the hardest thinking to AI erodes the very capacities it was meant to build. Learning involves struggle. We should be careful, though, not to conflate the effects of AI technologies used that way with those of a learning management system that, thanks to AI technologies, supports a more proactive, dynamic, and personalized learning experience.

Next, there is the surveillance a platform that observes learning requires — the constant observation, recording, and projection of a student's process, and the freedom to fail in private that it puts at risk.

There is platform concentration and operational fragility: when a single platform sits at the center of teaching, learning, and communication, an outage at the vendor is an outage of academic continuity, and a breach is everyone's breach — as the Canvas failure that opens this essay, and the one that followed it, both showed. And there is the risk of a more pedagogically active LMS built and controlled by a small number of well-capitalized vendors whose commercial interests are not necessarily those of students or schools. The field has been burned before — Knewton raised more than $180 million on its "robot tutor in the sky" and hyper-personalized learning, and was sold in a fire sale a decade later without

delivering — so the next round of personalization promises deserves skepticism.

There is also the temptation to dress commercial advantage as caution. Academic integrity concerns are real — institutions need confidence that the entity engaging with the learning environment is the enrolled human, not software acting on her behalf — but those same concerns can become cover for a major LMS player to lock down its APIs, charge tolls for partner status, and narrow who gets to integrate at all, in the name of safety. We must guard against opportunism dressed in the language of integrity, especially as the balance between platform-level and ecosystem-level services shifts.

Whether AI technologies help or harm is a function of how they are deployed and governed — who decides what the platform does, with what data, and under whose policies. That is the question every one of these risks finally points to.

I do not debate each risk here because they are not inevitable, even as they are most definitely possible. The pragmatist in me believes it ultimately comes down to execution. If you are wary of the essay that names a risk only to wave it away in the same breath, fair enough. But these risks are not settled in abstract debate; they are settled in the building process.

In university architecture and campus planning there is something called the desire path: rather than pre-designing every walkway around a new building, planners leave the grass intact, watch where people actually wear a trail, and pave that. A learning platform develops in the same way — with intentional design and in the paths faculty and students wear into it — and so do its failures. You see where tracking becomes surveillance, where the offloading sets in, where the lock-in tempts, and you respond.

As incoming CEO of Blackboard, I will be in a position to shape how at least one player executes — and to watch those paths form up close. While this essay is decidedly optimistic, with a strong point of view about the future of digital learning infrastructure in higher education, I am a "radical incrementalist." Change in education rarely takes shape with A leading directly to D; A leads to B, which leads to C, and then D. Radical ends, incremental means — one benefit of which is that it lets you calibrate progress against these concerns, rather than standing still or betting the whole design on a blueprint drawn from theory.

In a recent report on leading faculty through this moment, Jeff Selingo argues that what stalls AI in higher education is seldom a matter of tools or training; the real barriers, he writes, are "cultural and psychological rather than technical." What looks like faculty resistance, in his telling, usually runs deeper — an unease about whether they still command their craft, control how they teach, and have a place in an institution that is changing under them. A platform can make new approaches practical; it cannot supply the stability, the hands-on experience, or the trust that decide whether faculty adopt them. That is another reason to move incrementally — you earn adoption in low-stakes experiments and small wins, not by decree. This is what helped Blackboard gain broad adoption by faculty on campuses in the early days of the Internet.

In fairness to Hopkins, he recognized a core truth. He was wrong to resist the apparatus — Williams needed its libraries and labs, as learning needs its digital platforms — but right that the apparatus is never the teaching. It can serve the instructional relationship, but it cannot replace it. Education is a social process. AI technologies will not decide on their own whether this moment strengthens or weakens the human core of teaching and learning. They are tools; ordinary technologies. We will decide.

Coda: The Log

Which returns us to the “radical end” of a transformed digital log. When we founded Blackboard, I wrote that the best missions are the ones never quite achievable, the kind that keep you busy for a lifetime; our vision was to make the Internet a powerful environment for learning, and I quoted Browning – a man’s reach should exceed his grasp. We named the company after one of the oldest platform for teaching there is. Thirty years on, the platform can finally organize itself around the learner, carry the work of learning, and hold a record of what was learned – earning, in fact and not just in name, the L it was later given. If the credential has long certified that a student sat at the other end of the log — seat time, and a grade above C, not what she learned — the platform can finally enable it to represent (and motivate) the learning itself. The builders were right at Williams. They are right now. We can do better than a log.

Matthew Pittinsky is the incoming CEO of Blackboard, which he co-founded in 1997 and previously helped lead as CEO and Executive Chairman. More at matthewpittinsky.com.

AUTHOR’S NOTE: *Mark Hopkins's Log* draws on and synthesizes a cluster of working drafts I wrote while developing several discreet, but interrelated arguments. For readers who want to dig deeper, visit matthewpittinsky.com to read the working drafts. (The essay includes verbatim text I originally wrote in the working drafts.) As well, since I am writing about AI I decided to use AI in my writing process. While nothing I do is particularly novel, if you are interested in how I use AI to write I've drafted this note.