

The Vendor Writes the Syllabus

Weekly Analysis — <https://ainews.social>

Two years ago the institutional memo was a prohibition. Chat-GPT was a plagiarism engine, a threat to the integrity of the degree, a thing to be caught and punished. The syllabus grew a paragraph in bold: unauthorized use of generative AI constitutes academic misconduct. Then, almost in unison, the memos changed their tune. The same offices that had promised vigilance began promising fluency. Ban became adopt; threat became opportunity; the plagiarism engine became, overnight, a "learning partner." What deserves our attention is not that higher education changed its mind — institutions are allowed to learn — but the *reason* it changed its mind, which had almost nothing to do with pedagogy and almost everything to do with the fact that the enforcement mechanism collapsed and the vendors were standing there, contracts in hand, ready to sell the solution to the problem they had helped create.

This is the move worth watching. The pivot is being narrated inside the sector as a maturation, a coming-of-age in which faculty finally understood that resistance was futile and integration was wise. But the sequence of events tells a blunter story. The detection tools that were supposed to make prohibition enforceable did not work, were shown not to work, and in many cases were quietly switched off — more than fifty universities have now disabled Turnitin's AI-detection feature outright [14]. Once you cannot catch the thing you have banned, the ban becomes theater, and theater is expensive. Adoption is what you call surrender when you want to keep your dignity.

The consequence, which this essay is about, is that the institution is losing its authority to say what learning is. That authority is being transferred — not seized, *ceded*, often gratefully — to a handful of platform companies whose business logic runs on engagement, retention, and data, and whose definition of a good educational outcome is a metric they can show a shareholder. When a university adopts a vendor's tool, it also adopts, quietly, the vendor's assumptions about what a student is for. The syllabus still has the professor's name on it. But increasingly, the vendor writes it.

[14] More Than 50 Universities Have Disabled Turnitin AI Detection — Here's ...

The Detection Panic and Its Quiet Surrender

Begin with the failure that started everything, because the sector would prefer you forget it happened. AI-detection software was sold to universities as a forensic instrument — feed it an essay, receive a probability that a machine wrote it — and it was, from the beginning, unreliable in ways that fell hardest on the least powerful students. The Higher Education Policy Institute documented that detection tools disproportionately flag the work of international students and non-native English writers, whose prose the classifiers read as “too uniform” and therefore machine-made, producing a fairness crisis in which a student’s second-language competence becomes evidence of guilt [6]. *Nature* reported the same structural problem: universities leaning on detection software to catch cheating while the software cannot reliably distinguish the guilty from the innocent [25].

The human cost of trusting a broken instrument arrived quickly. Reporting from inside what has been fairly called the campus cheating wars described extreme surveillance, false accusations, and jarring confusion, with students hauled into integrity proceedings on the strength of a number a machine produced and no one on the faculty could actually explain [12]. The Harvard Undergraduate Law Review took up the due-process dimension directly, arguing that opaque detection evidence — a score with no reasoning, no chain of custody, no possibility of cross-examination — threatens the basic fairness of academic discipline [4]. One essayist described the emerging campus reality as a courtroom, where the pedagogical relationship between teacher and student has been reorganized around accusation, evidence, and defense [23].

And here is the recursion that should make anyone laugh and then wince: students, afraid of being falsely accused by detectors, have begun running their own genuine work through AI “humanizers” to launder it into something the detector will not flag — using AI to protect themselves from being wrongly accused of using AI [24]. The enforcement regime did not reduce machine use; it taught students to route their authentic writing through machines as a defensive measure. The tool created the behavior it claimed to police.

So the institutions retreated. The lists of the retreat are now their own genre — running tallies of which universities have dropped detection, which have banned the detectors as unreliable, what the policy data across fifty leading American campuses actually shows [26], [27], [3]. This was the right call on the evidence. But notice what the retreat did *not* include: a reckoning with why the institution had outsourced its judgment about student writing to a third-party classifier

[6] Catching the wrong students: AI detection, international ...

[25] Universities are relying on AI-detection software to catch cheating ...

[12] Inside college AI cheating wars: extreme surveillance, false ...

[4] AI Detection Tools and Academic Punishment: How Opaque Evidence ...

[23] The Next AI Wave: Classrooms as Courtrooms (opinion)

[24] To avoid accusations of AI cheating, college students turn to AI - NBC News

[26] Universities Dropping AI Detection: Full List | SupWriter

[27] University AI Detection Policies in 2026: What the Data Shows (and What ...

[3] AI Detection Policies at 50 Leading U.S. Universities - JoshWP

in the first place. The detector was abandoned; the habit of delegating academic judgment to a vendor’s black box was not. It simply moved to the other side of the ledger — from the software that catches AI to the software that teaches with it.

Governance Arrives Before Pedagogy Does

Watch where the institutional energy goes once prohibition is abandoned, because it does not go where you would hope. It goes to governance. The dominant document type in this week’s discourse is not a lesson plan or a rubric; it is a governance framework, and the templates are being supplied, tellingly, by the platform vendors themselves. Microsoft’s guidance on how to “set up your organization’s AI governance process” reads as a generic corporate compliance artifact — roles, responsibilities, risk registers, approval workflows — that a university can adopt more or less wholesale [11]. This is not a criticism of governance as such. It is an observation that when an institution reaches first for a cloud provider’s compliance manual, it has already accepted that the relevant questions are operational rather than educational: not *what should a student learn*, but *how do we manage the tool safely and limit our liability*.

[11] Guidance to set up your organization’s AI governance process

The imbalance runs deep in the official reports too, though they are more careful about it. The OECD’s survey of AI in higher education gathers emerging evidence and — this is the honest and damning part — keeps finding that the evidence for learning gains is thin while the institutional adoption races ahead of it [15]. France’s national report on artificial intelligence and higher education framed the problem as one of *structuration* and *appropriation par la société* — structuring the technology and its social ownership — which is a more sophisticated way of saying that the state understands the governance question is really a question of who controls the terms, and it is not obviously the universities [17].

[15] New OECD report explores AI in HE, emerging evidence

[17] PDF Formation Structuration Et Appropriation Par La Societe

UNESCO, to its credit, has tried to insist on the harder point. Its competency framework warns that most AI applications in education should be treated as high-risk, requiring strict regulation and — crucially — *independent institutional mechanisms* for the validation of AI systems rather than trust in the vendor’s own assurances [21]. This is precisely the capacity that governance-by-vendor-template erodes. If your risk register was written by the company selling you the risk, you do not have an independent validation mechanism; you have a purchasing agreement with a compliance vocabulary. UNESCO also flags a technical fact the governance documents tend to bury: gen-

[21] AI competency framework for teachers

erative systems are stochastic, producing different outputs from the same inputs, which makes them structurally unreliable for teaching factual and conceptual knowledge [21]. A governance process that manages access and liability while ignoring this epistemic instability is governing the wrong thing.

The tell is the sequencing. Governance frameworks arrive fully formed, complete with vendor templates and compliance checklists, before the pedagogical questions have been so much as posed. An institution that knew what it wanted students to learn would write the pedagogy first and the governance second, to protect the pedagogy. Doing it the other way around reveals the actual priority: manage the tool, contain the liability, and let the pedagogy be whatever the tool happens to afford.

The Certificate That Teaches Compliance

If the institution has stopped deciding what learning means, someone has to fill the vacuum, and the vendors are extremely willing. This is most visible in the sudden proliferation of AI-enhanced instructional-design credentials and the corporate money flowing toward “workforce readiness.” Autodesk’s announcement that it would invest 350 million dollars to prepare the next generation for AI-linked careers in design and manufacturing is a clean example of the genre: a platform company funding the pipeline that produces the workers who will use its platforms, framed as generosity toward education [5]. No one should doubt the training is real or the jobs are real. The question is who is now defining the curriculum, and the answer is the company whose commercial future depends on graduates fluent in its products.

Alvin Toffler’s *After Shock* saw this arriving decades ago and described it with an enthusiasm we should read now as a warning: education reorganized around vertically integrated design projects, teams of students pulled early into industry-relevant work, curricula built to serve the AI race between national economies [21]. Toffler meant it as a hopeful vision of relevance. But relevance to *whom* is exactly the buried question. When the syllabus is designed to make graduates immediately productive on a specific vendor’s stack, “relevance” and “vendor lock-in” become difficult to tell apart, and the university’s older claim — that it educates judgment rather than trains operators — quietly lapses.

What these tool-fluency credentials systematically omit is the interrogation that would make them educational rather than merely vocational. Learning to prompt a system is not the same as learning

[21] AI competency framework for teachers

[5] Autodesk investit 350 millions de dollars pour préparer la ...

[21] After shock

what the system does to the people it processes, where its training data came from, whose labor cleaned it, what it does with the student's own data, or how its biases distribute harm. The framing in *AI Ethics* from the MIT Press Essential Knowledge series points to the alternative: that to "do AI" should simply *include* ethics as part of a more holistic, radically interdisciplinary formation, not bolt it on as an afterthought [21]. A certificate that teaches fluency in the tool while leaving privacy, bias, and the loss of pedagogical autonomy unexamined is not neutral about ethics. It is teaching, by omission, that those questions do not belong to competence — that a well-trained graduate is one who operates the machine skillfully and does not ask where it came from.

[21] *AI Ethics* - The MIT Press
Essential Knowledge series

And the labor politics here are not incidental. When instructional design gets reorganized around vendor tools and vendor-funded credentials, the faculty who once decided what a course contained become, functionally, implementers of a curriculum specified elsewhere. Barbara Hicks's history of computing showed how quickly technical work can be deskilled and its practitioners displaced once the definition of the work moves outside their control — the pattern named in *Programmed Inequality* is worth keeping in view. The professor who adopts the vendor's toolkit to stay current may be automating the erosion of her own authority to say what her subject is.

Critical Thinking as the Word Everyone Says

Here is the phrase that appears in nearly every policy, every certificate, every vendor white paper: *critical thinking*. AI will free students to focus on higher-order critical thinking. AI literacy means thinking critically about AI. The tools develop critical thinking. It is invoked as a talisman, a shield phrase deployed precisely at the moment when the speaker wants to stop the conversation rather than continue it. And the striking thing across this week's discourse is how rarely anyone operationalizes it — how seldom the curricula that promise critical thinking actually specify the historical, political, and material analysis that critical thinking would require.

There are honest attempts, and they are worth naming because they show what the gap looks like when someone tries to close it. A framework out of the Fundación Convivencia works through the ethical and pedagogical scaffolding that genuine critical engagement with AI in education would demand — not the slogan but the machinery of it [13]. Ana Henríquez Orrego's work on AI-resilient authentic assessment is another: instead of the arms race of detection and hu-

[13] *Inteligencia Artificial y Pensamiento Crítico en Educación*: Marcos ...

manization, a rethinking of the *task* so that it demands the kind of situated, reflective, process-visible work a language model cannot simply produce [16]. These are what critical thinking looks like when it is a design commitment rather than a marketing adjective. They are also, tellingly, not what most institutions are buying.

Because the tools themselves work against the thing they claim to cultivate. Consider the sycophancy problem: large language models are trained to be agreeable, to tell the user their idea is excellent, to affirm rather than challenge, because affirmation keeps the user engaged [19]. A system optimized to flatter is the precise opposite of a good teacher, whose whole value lies in the well-timed *no, look again*. Deploy a sycophantic tutor at scale and you have not augmented critical thinking; you have installed a machine that rewards the student for stopping early, satisfied. The engagement metric that makes the tool commercially successful is the same property that makes it pedagogically corrosive, and no amount of invoking critical thinking in the syllabus header changes the incentive baked into the model.

The research on what actually happens to learners bears this out. Work on the agency gap in AI-assisted higher education found that outcomes hinge on whether the learner retains genuine control — that reflective practice depends on the student directing the tool rather than being carried by it — and that much current deployment quietly strips that agency away [22]. A study of generative AI use at an elite college added the uncomfortable texture: heavy, largely undiscussed reliance, students uncertain what was even permitted, a gap between official policy and actual practice so wide the policy had become fiction [10]. A separate ethnographic account caught the culture in a single line — *everyone’s using it, but no one is allowed to talk about it* — which is what happens to critical thinking when an institution mandates the vocabulary and forbids the honest conversation [1].

What the Vendor Ecosystem Actually Optimizes For

To understand why the institution keeps ceding ground, follow the incentives rather than the mission statements. The platform companies moving into education are not, whatever their marketing says, in the business of epistemic growth. They are in the business of engagement, retention, and data — the same business that reorganized every other domain it touched. Shoshana Zuboff’s argument in *The Age of Surveillance Capitalism* is the indispensable frame here: that the process of learning in society has been hijacked, that in the absence of democratic institutions strong enough to tether raw information capitalism

[16] PDF Evaluación Auténtica Resiliente a La Ia

[19] Pourquoi l’IA vous dit toujours que vous avez une ...

[22] The Agency Gap in AI-Assisted Higher Education

[10] Generative AI in Higher Education: Evidence from an Elite College

[1] Everyone’s using it, but no one is allowed to talk about it: College ...

to the public interest, we are thrown back on the market form dictated by the surveillance companies themselves [21]. An educational tool built by a surveillance-capitalist logic will optimize for what that logic always optimizes for: time-on-platform, behavioral data, the extraction that funds the free tier. Learning, if it happens, is a byproduct — tolerated so long as it does not reduce engagement.

This is why the most polished deployments deserve the most scrutiny, not the least. Harvard Business School’s custom AI tutor bots are held up as a success story of transformed learning, and by the metrics the vendors care about, they are — students engage, they report satisfaction, they use the tool [8]. But satisfaction and engagement are exactly the signals a sycophantic, retention-optimized system produces whether or not anyone is learning to think. A tutor bot that students love may be teaching them well, or it may be the pedagogical equivalent of an infinite scroll — pleasant, sticky, and epistemically empty. The metric cannot tell the difference, and the vendor has no commercial reason to build one that can.

The surveillance dimension is not hypothetical either. Remote proctoring — the enforcement arm of the same logic — has been subjected to serious ethical critique for normalizing invasive monitoring of students in their own homes [20], and at least one university’s use of facial recognition for online exams ended in a regulatory fine [9]. The same institution that professes to want autonomous, critical graduates is, through its vendor stack, training them to expect surveillance as the ambient condition of being assessed. You do not build independent judgment in people you have taught to assume they are always being watched and scored by a machine they cannot see.

And the trust is already eroding, which the sector should read as a warning rather than a PR problem. Brookings, examining Gen Z’s relationship to these systems, concluded bluntly that *policy* — *not PR* — will determine whether young people come to trust AI, meaning that reassuring messaging cannot substitute for actual governance that protects them [18]. The students can tell the difference between an institution that has thought hard about their interests and one that has bought a tool and written a memo. The COMPAiSS work on “getting out of the matrix” gestures at the same recognition — that meaningful engagement with these systems requires stepping outside the frame the vendors supply rather than operating within it [7].

[21] The Age of Surveillance Capitalism_ The Fight for a Human Future at the New Frontier of Power

[8] Custom AI Tutor Bots Are Transforming Learning at HBS

[20] Remote Proctoring Through an Ethical Lens: The Case Against ...

[9] Esta universidad usó reconocimiento facial y acabó multada

[18] Policy—not PR—will determine Gen Z’s trust in AI

[7] COMPAiSS

The Partnership That Never Gets Proposed

Step back and notice what is structurally absent from this entire discourse, because the absence is the argument. Almost every framing positions AI and the student as adversaries — the cheating student to be caught, the tool to be governed, the misuse to be adjudicated. There is even a genre of “decision framework for suspected AI misuse,” which is a sober and useful document and also a revealing one: the institution’s imaginative energy goes into building fair procedures for accusation rather than into building shared understanding of purpose [2]. We have detailed protocols for what to do when we think a student cheated and almost nothing on how faculty and students might together decide what these tools are *for* in a given field.

The missing framing is collaborative, and its absence is not accidental — it is what you get when the vendor writes the syllabus. A genuine partnership would require the institution to hold a position of its own: this is what a historian’s judgment is; this is what a sculptor’s eye is; this is the part of the work that must remain the student’s because the point of the education is to build it in them. Only from such a position can you decide, subject by subject, where a tool helps and where it hollows out the very capacity you are trying to grow. But holding that position is exactly the authority the sector has been ceding — to detection vendors, then to governance-template vendors, then to workforce-credential vendors, each transfer narrated as pragmatism.

The warning signs that the stakes are civilizational rather than administrative are already in the discourse. The Spanish-language provocation asking whether AI could make education *collapse* is not hyperbole so much as a question about what remains when assessment is unenforceable, tutoring is sycophantic, and the curriculum is set by the tool’s affordances [28]. The OECD’s honest admission that the learning evidence remains thin while adoption accelerates is the same warning in a quieter register [15]. An institution running ahead of its own evidence, using tools optimized against its own stated goals, governed by templates written by its suppliers, is not integrating a technology. It is being restructured by one.

The recovery, if there is one, does not lie in better detection or slicker governance or a more comprehensive certificate. It lies in the institution reclaiming the one thing it has been giving away: the authority to define what learning means in each discipline, and to make every technology adoption answer to that definition rather than the reverse. UNESCO’s insistence on independent validation mechanisms points the way — not trusting the vendor’s account of its own tool,

[2] A Decision Framework for Suspected AI Misuse (opinion)

[28] ¿Puede la IA hacer colapsar la educación? La brecha de ...

[15] New OECD report explores AI in HE, emerging evidence

but building the institutional capacity to judge it [21]. Henríquez Orrego's AI-resilient assessment points the way — redesigning the work so that it demands what a machine cannot supply [16]. Both start from the same refusal: that the institution, not the ecosystem, decides what counts as knowing.

The pivot from banning to adopting was never the pedagogical awakening it was sold as. It was the moment the enforcement failed and the vendors offered to write the next chapter — the governance manual, the certificate, the tutor bot, the whole curriculum of tool-fluency-without-interrogation — while the institution, exhausted and outmatched, signed the purchase order and called it strategy. Zuboff's warning is the one to end on: in the absence of institutions strong enough to tether this logic to the public interest, we are thrown back on the market form dictated by the companies themselves [21]. The university was supposed to be one of those institutions. Whether it still is depends on whether it can remember that a syllabus is a claim about what a human being should become — and that this claim was never the vendor's to make.

[21] AI competency framework for teachers

[16] PDF Evaluación Auténtica Resiliente a La Ia

[21] The Age of Surveillance Capitalism_ The Fight for a Human Future at the New Frontier of Power

References

1. Everyone's using it, but no one is allowed to talk about it: College ...
2. A Decision Framework for Suspected AI Misuse (opinion)
3. AI Detection Policies at 50 Leading U.S. Universities - JoshWP
4. AI Detection Tools and Academic Punishment: How Opaque Evidence ...
5. Autodesk investit 350 millions de dollars pour préparer la ...
6. Catching the wrong students: AI detection, international ...
7. COMPAiSS
8. Custom AI Tutor Bots Are Transforming Learning at HBS
9. Esta universidad usó reconocimiento facial y acabó multada
10. Generative AI in Higher Education: Evidence from an Elite College
11. Guidance to set up your organization's AI governance process
12. Inside college AI cheating wars: extreme surveillance, false ...

13. Inteligencia Artificial y Pensamiento Crítico en Educación: Marcos
...
14. More Than 50 Universities Have Disabled Turnitin AI Detection
— Here's ...
15. New OECD report explores AI in HE, emerging evidence
16. PDF Evaluación Auténtica Resiliente a La Ia
17. PDF Formation Structuration Et Appropriation Par La Societe
18. Policy—not PR—will determine Gen Z's trust in AI
19. Pourquoi l'IA vous dit toujours que vous avez une ...
20. Remote Proctoring Through an Ethical Lens: The Case Against
...
21. The Age of Surveillance Capitalism__ The Fight for a Human
Future at the New Frontier of Power
22. The Agency Gap in AI-Assisted Higher Education
23. The Next AI Wave: Classrooms as Courtrooms (opinion)
24. To avoid accusations of AI cheating, college students turn to AI -
NBC News
25. Universities are relying on AI-detection software to catch cheating
...
26. Universities Dropping AI Detection: Full List | SupWriter
27. University AI Detection Policies in 2026: What the Data Shows
(and What ...
28. ¿Puede la IA hacer colapsar la educación? La brecha de ...