The Accommodation Trap: When AI Widens the Gap It Was Meant to Close

Weekly Analysis — https://ainews.social

There is a particular kind of institutional failure that announces itself as success. A lecture-capture system is rolled out across a university; the press release celebrates accessibility; the disability services office is told its budget can be redirected; and three months later a deaf graduate student in engineering discovers that the automatic captions, when her advisor speaks — he is Nigerian — render "stochastic gradient" as "stick acid grading" and her notes are useless. The captions are working. They are not working for her. The institution will, if pressed, call this an edge case. It is not an edge case. It is the structural signature of how artificial intelligence is arriving in higher education: tools that average well across a population while failing catastrophically along its margins, and an institution that has learned to measure averages.

This essay is about that pattern, and about how a sector that has spent a decade talking about equity is on the verge of automating its inequities at scale. The discourse this week — across faculty senate minutes, vendor white papers, accreditor guidance, and an unusually candid Delphi study of 70 international experts [13] — converges on a single fact that the press releases have not absorbed: AI deployment in higher education is producing a double-edged outcome, and the second edge cuts the students the institution claims to serve most.

[13] Governing generative AI in higher education: a global Delphi ...

The pattern is now well enough documented that "we didn't know" is no longer available as a defense. What is at stake is whether the institution treats AI as a procurement question — sign the contract, pilot the tool, count the logins — or as a governance question that begins where procurement ends. The argument here is that the procurement frame is itself the problem, because it locates AI's harms outside the academic mission rather than inside it, and because it consistently assigns the cost of failure to the people with the least power to refuse.

The Two-Tier Campus Hidden Inside the Glossy Pilot

The cleanest way to see the accommodation trap is to look at what happens when an AI tool is described as universally beneficial. Automatic captioning is the canonical case. For students with hearing impairments, the technology is genuinely transformative; nobody serious disputes this. But the same systems that handle a mid-Atlantic accent in a quiet seminar room perform measurably worse on instructors with non-English first languages, on Black English vernacular, on lectures held in rooms with HVAC noise, on any course where the instructor moves and the lavalier microphone rustles. The student who needs the caption most — the one whose own English is shaky, whose professor's English is shaky, whose Wi-Fi drops in the back of the lecture hall — gets the worst caption. Karen Lumsden's audit work has pointed out that 41% of UK universities have no publicly available AI policy at all [17], which means that in nearly half the sector there is no document a student could even point to in order to complain.

[17] Karen Lumsden, PhD's Post

What the captioning case reveals at small scale, infrastructure reveals at large. The discourse this year has begun to acknowledge — quietly, almost embarrassed — that the gleaming pilot in the flagship classroom rests on a substrate that is not equally distributed. A recent overview of generative AI in higher education makes the point that access to reliable bandwidth, current devices, and the paid tiers of tools like ChatGPT Plus or Claude Pro is now a material factor in academic performance, not a peripheral concern [11]. On the campuses where 70% of students live off-campus and a third work more than twenty hours a week, "AI literacy" is being measured against students whose home Wi-Fi cannot stream a tutoring session without buffering. The institution treats this as a personal misfortune. It is in fact the institution's procurement decision arriving at the student's apartment.

[11] Generative AI in Higher Education

The Indian case sharpens the point. OpenAI's expansion into Indian higher education through institutional partnerships [18] is framed as democratization, and at one level it is — students at partner universities gain free or subsidized access to tools that cost real money elsewhere. But the partnership is also a logistical fact about which Indian universities have the bandwidth, the IT staff, and the English-language administrative posture to make the deal in the first place. The hundred elite institutions that sign on become the showcase. The two thousand that don't become the control group in an experiment nobody designed and nobody is measuring.

[18] OpenAI s'étend dans l'enseignement supérieur indien via des ...

Advising Algorithms and the Optimization of Attrition

If captioning is the visible case of the accommodation trap, AI advising is the invisible one. The retention-analytics industry has been selling colleges a story for over a decade: deploy our model, and we will identify struggling students earlier, intervene more precisely, and lift your graduation rates. The story is now substantiated by enough data that the headline claim is plausibly true — on average. The trouble is what "on average" hides. Recent policy analysis describes the rise of what its authors call the "algorithmic institution," in which artificial intelligence is deployed less as a pedagogical tool than as a policy response to a sector in financial and demographic crisis [25]. The model is optimizing for institutional survival metrics — credit hours completed, time-to-degree, retention rate — that align imperfectly with student flourishing and sometimes oppose it.

[25] Risk, Retention, and the Algorithmic Institution: Artificial Intelligence as a Policy Response to Higher Education in Crisis

The mechanism that produces the inequity is worth describing in detail because it is so frequently misunderstood. An advising algorithm trained on historical outcomes learns that first-generation students from low-income high schools have, in the past, withdrawn from organic chemistry at higher rates. The algorithm flags such a student in week three, when she gets a 68 on the first midterm. An advisor reaches out, suggests she drop to a lower-track section, perhaps shift majors. The intervention is offered with care; the advisor genuinely believes she is helping. The student, who was on track to recover — who had, in fact, already booked office hours — interprets the message as the institution telling her she does not belong. She switches majors. The algorithm records a successful retention. The pre-med pipeline loses another first-generation student. The model's accuracy improves.

This is not a hypothetical. The pattern of over-surveillance triggering counterproductive interventions has been documented across several K-12 deployments of tools like Gaggle, GoGuardian, and Bark, where automated flagging has produced false alarms, disciplinary actions, and in some cases arrests of students whose searches or essays were misread by the model [26]. The Spanish-language coverage of the same systems makes the point even more bluntly: the false alarms have led to punishments and arrests that fall disproportionately on the students whose behavior the algorithm was least equipped to read [8]. Higher education has imported the same surveillance posture under gentler names — "early alert," "student success platform," "predictive advising" — and it is producing, predictably, the same pattern at a more sophisticated layer of harm.

[26] School AI surveillance like Gaggle can lead to false alarms, arrests ...

[8] Falsas alarmas de vigilancia con IA han provocan castigos y arrestos ...

Ruha Benjamin's argument in [3] is the right frame here. The book's point is not that algorithms are racist in some crude sense, but

[3] Race After Technology

that systems built on historical data inherit the discriminatory shape of the past and re-export it as neutral prediction — what she calls the New Jim Code. When the institution adopts a retention algorithm, it is not adopting a tool; it is adopting a theory of which students are worth keeping, encoded in the training data, and then washing that theory in the disinfectant of mathematics. The administrator who deploys the system can no longer be accused of bias because the bias is now the model's, and the model is "just" predicting.

The Plagiarism Panic and the Pedagogical Vacuum It Exposed

The other half of this year's higher-education AI discourse has been about cheating, and it has been remarkably ugly. Adelphi University spent much of the spring defending itself against a lawsuit from a student it accused of submitting AI-generated work on the basis of a detector's score [2]. French jurists have begun to ask, with some asperity, whether a university can sanction a student for AI use when the university itself has produced no written rule about it [16]. Canadian survey data shows that roughly one student in three admits to transgressing institutional rules with AI [31], a number that has the additional effect of suggesting that whatever the institution thinks its rules are, the students have not received them with the clarity the institution imagines.

[2] Adelphi University accused a student of using AI to ... - Newsday

[16] Intelligence artificielle : l'université peut-elle sanctionner sans règle

[31] Un étudiant sur 3 transgresse les règles à l'aide de l'IA

What is striking about the cheating panic — and this is where the accommodation frame returns — is who pays its costs. The students most likely to be falsely accused are the ones whose writing the detector has the most trouble reading: non-native English speakers, students whose prose is unusually formal because they learned academic writing from a tutor rather than from immersion, neurodivergent students whose syntax is idiosyncratic. The Adelphi case is one visible instance; the invisible cases are the students who, having seen what happened, simply stop writing in their own voice and start writing in the voice the detector expects, which is to say, they let the detector teach them to write like a confident undergraduate from an American suburb. The accommodation trap, again: a tool aimed at policing one population teaches another to dissimulate.

Harvard's coverage of how to preserve learning amid AI shortcuts gets at the deeper pedagogical question [21] — namely, that the moment a faculty member chooses to police output rather than redesign the assignment, the institution has surrendered its pedagogical authority to the vendor. The detector becomes the curriculum. The argument for friction — that learning requires productive resistance,

[21] Preserving learning in the age of AI shortcuts — Harvard Gazette

that frictionless tools can hollow out the cognitive work an assignment is supposed to occasion — has been made carefully in a recent piece on AI-mediated information seeking [27], and it deserves to be taken seriously by anyone designing a syllabus. The friction argument is not nostalgia. It is the insistence that an assignment whose only barrier to completion is whether the student can paste a prompt is not, in any meaningful sense, an assignment.

[27] The case for friction in AI-mediated information seeking and learning

The empirical record on this is more ambiguous than either the boosters or the panickers will admit. There is evidence that generative AI reduced study time on math problems without compromising performance in some controlled settings [12], which sounds like a clean win until one asks what the saved time was used for — additional learning, or simply less learning. A French overview of the impact of generative AI on critical thinking is more cautious, noting consistent signals that frequent AI use is associated with reduced engagement in the cognitive steps that critical reasoning actually requires [14]. A separate piece on reading, critical thinking, and problem-solving among students using AI tools finds the effects bifurcated: students who already had strong metacognitive habits used AI to deepen them, while students who lacked those habits used AI to substitute for them [28]. The accommodation trap, restated in cognitive terms: the tool helps the prepared and harms the unprepared, and the institution is celebrating the average.

[12] Generative AI Reduced Study Time on Math Problems and ...

[14] Impact de l'IA générative sur la « pensée critique »

[28] The Impact of AI on Students' Reading, Critical Thinking, and Problem ...

When Personalization Works — and the Case It Cannot Be Generalized

It would be dishonest to write this essay without naming the cases where AI in higher education has actually delivered on its promises, because those cases matter and because they are the cases the institution is, understandably, drawing the wrong lessons from. The most rigorous of these is the Harvard physics tutor study, in which a faculty member designed a course-specific AI tutor calibrated to the curriculum and pedagogy of an introductory physics course; engagement roughly doubled and learning outcomes improved [22]. A study of an "AI Digital Teacher" in a human-AI collaborative learning configuration found similarly promising results when the tool was used as a structured collaborator rather than a substitute [29]. Engagement metrics in broader survey work on student use of AI tools show that meaningful engagement is achievable when the design matches the learning task [9].

[22] Professor tailored AI tutor to physics course. Engagement doubled.

[29] The impact of an AI Digital Teacher on human-AI collaborative learning in higher education

[9] Frontiers | Student engagement with AI tools in learning: evidence from ...

What unites these successes is what is missing from the procurement playbook the rest of the sector is running. The Harvard physics tutor worked because a faculty member with disciplinary expertise spent significant time tailoring the tool to the specific cognitive moves the course was trying to occasion. The successful Digital Teacher deployments were ones designed around collaboration rather than replacement. The studies that show engagement gains are studies of tools embedded in pedagogy, not bolted onto it. The lesson the institution wants to draw — ”AI tutors work, so let’s license one for every course” — is the opposite of the lesson the evidence supports, which is ”AI tools work when faculty have the time, expertise, and authority to design them into a course.” The first lesson takes a procurement contract. The second takes a labor settlement.

This is the place to mention what the discourse this year has begun to call the graduate-teaching-labor problem. A pointed working paper on AI and graduate teaching labor argues that AI tools are reshaping workload, autonomy, and the structure of academic work in ways that fall hardest on the contingent workforce — graduate students, adjuncts, postdocs — who do most of the actual teaching in research universities [19]. The tenured faculty member who builds a tailored AI tutor for her physics course is the celebrated case; the adjunct teaching four sections at three institutions who is told by an administrator to ”use AI to grade faster” is the unmentioned case, and the unmentioned case is the more common one by orders of magnitude. AI in higher education is not, primarily, a tool being deployed by faculty. It is a tool being deployed at faculty, and the faculty most likely to be on the receiving end of that deployment are the ones with the least standing to refuse.

[19] PDF AI and Graduate Teaching Labor: Reshaping Workload, Autonomy, and ...

The Governance Vacuum and Who Is Filling It

Into the space where institutional policy ought to be, three actors have moved. The first is the vendor, who has both the urgency and the resources to fill a void. The second is the individual faculty member, improvising syllabus by syllabus. The third is the student, who in the absence of clear rules is making decisions about disclosure, attribution, and academic risk on the basis of rumor. None of these actors is the right one to be setting institutional AI policy, and the fact that they are is the governance failure the sector has not yet named clearly.

The international Delphi study of governance frameworks for generative AI in higher education is unusually clear-eyed about this [13]. Its experts converge on a small number of principles — inclusive design, transparent assessment, redress mechanisms, ongoing audit — that are

[13] Governing generative AI in higher education: a global Delphi ...

notable mostly for how completely absent they are from the procurement contracts most universities have signed. A separate course-design proposal on AI governance in higher education argues that universities need not merely to write AI policies but to teach AI governance as a curricular subject, because the regulatory environment is now developing faster than any single policy document can track [5]. The Quebec student-perspective document on AI systems makes a related point from the other direction: students themselves have begun to articulate, in remarkably sophisticated terms, what fair institutional AI use looks like, and they are not being listened to [20].

[5] AI Governance in Higher Education: A course design exploring regulatory ...

[20] PDF Perspective Étudiante Sur Les Systèmes D'Intellige Nce Artificielle ...

The trouble with the governance discourse, as it currently stands, is that it has been almost entirely procedural — what committees should exist, what disclosures should be required, what risk tiers should be assigned. These are real questions and the answers matter, but the procedural frame obscures the more difficult substantive question, which is what the institution is actually for. The Spanish-language educational scholarship on artificial intelligence and critical thinking in education has been pushing on this point, arguing that ethical frameworks for AI in education are downstream of a prior commitment about what education is meant to do [15]. A university that has decided its purpose is workforce throughput will write one AI policy; a university that has decided its purpose is the formation of critical citizens will write a different one. Most universities, having declined to decide, are writing both at once and acting on neither.

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

Kate Crawford's account in [3] is useful here, especially her point that AI systems do not merely process the world; they actively construct the categories — race, gender, ability, deservingness — through which the world becomes legible to institutions. When a university adopts an AI advising platform, it is not adopting a neutral measurement instrument. It is adopting a system that will, in the course of its operation, define what a "successful student" looks like, and it will define this in the terms the vendor's training data has set. The institution that buys this without examining the categories is not deploying a tool. It is outsourcing its anthropology.

[3] The Atlas of AI

Shoshana Zuboff's diagnosis in [3] sits behind this whole discussion. The vendor model that has colonized higher education is structurally identical to the model that colonized consumer attention: free or subsidized at the front end, monetized through behavioral data at the back end, with the institution serving as the unwilling but compliant collector. When a university signs a contract that gives a vendor access to keystroke-level student data in exchange for a discount on enterprise licenses, the university has accepted a deal whose terms it has not understood and whose costs its students will pay. The Delphi

[3] The Age of Surveillance Capitalism

experts' insistence on transparent assessment and redress mechanisms is a polite way of saying: the contracts most universities have signed do not contain either.

Faculty Are Moving — But Toward What?

One of this year's more interesting shifts is that faculty are no longer the bloc of refusal the boosters expected. A recent analysis of policy shifts in higher education documents faculty moving away from outright AI bans toward more nuanced engagement [7]. The shift is real and on balance probably healthy: a blanket ban was always going to be unenforceable and pedagogically unserious. But the move from "ban" to "engage" has not been accompanied by the institutional support — release time, training, redesign funding, technical assistance — that would make the engagement meaningful. The faculty member is being asked to redesign her course around AI on her own time, with her own judgment, against vendor tools she did not select, while the institution claims credit for the policy shift.

[7] Faculty Ditch AI Bans: Study Shows Policy Shift - AcademicJ…

This is where the LOGOS framework on human cognitive agency in AI-assisted assessment becomes useful as a diagnostic [30]. The framework's five levels distinguish between assessments where the AI is a tool the student uses, assessments where the AI is doing most of the cognitive work, and gradations in between. The point is not that any one level is correct, but that the institution that has not made a deliberate choice about which level it wants — and resourced that choice — is allowing the choice to be made for it, course by course, by the vendor's default settings. The competency-framework discourse in medical AI education makes the same point in a more applied register [1]: you cannot teach AI competently without first deciding what competence looks like, and most institutions have skipped that step.

[30] The LOGOS Framework: A Five-Level Taxonomy of Human Cognitive Agency in AI-Assisted Assessment

[1] A Competency Framework for Medical AI Education: Mixed Methods Study

UNESCO's [3] is worth invoking here precisely because it is so frequently invoked without being read. The framework's actual content is more demanding than the institutions that cite it would like to acknowledge: it asks teachers to develop critical capacity to assess both the positive and the negative impacts of AI tools, to recognize that generative systems are stochastic and therefore less trustworthy for factual and conceptual teaching, and to retain pedagogical authority over the curriculum rather than ceding it to the tool. The framework is not a permission slip. It is a job description that very few institutions have actually staffed.

[3] AI competency framework for teachers

The British public-attitude survey adds an exterior pressure that institutions have not yet metabolized: 20% of Britons think AI will produce civil unrest, and the public mood on AI and the future of work has shifted decisively toward fear over hope [23]. The university that brands itself as a leader in AI adoption is doing so in a public context where the public is increasingly skeptical of the technology and the institutions that are deploying it. The legitimacy question — touched on in a recent piece on GenAI in higher education, legitimacy, and laziness [10] — is going to get harder, not easier, and an institution whose AI strategy is ”be enthusiastic” will find itself increasingly out of step with the constituency whose tuition it depends on.

[23] Public have more fear than hope on AI and future of work, ...

[10] GenAI in Higher Education, Legitimacy and Laziness

What Co-Design Would Actually Require

The phrase that has been missing from this essay so far, deliberately, is ”inclusive co-design.” It is missing because it has become, in the institutional vernacular, a euphemism — the thing you say in the executive summary to indicate that you have thought about equity without committing to anything specific. Co-design is a concrete practice with concrete demands, and the institutions that have actually done it look measurably different from the ones that have not.

Concretely: co-design requires that the populations most affected by an AI deployment have authority — not consultation, authority — over its design, its rollout, and its discontinuation. It requires that the procurement contract include a clause permitting independent audit of the system's differential performance across demographic groups, and that the audit be funded by the institution and conducted by parties without a contractual relationship to the vendor. It requires a redress mechanism that a student can actually use, in the sense that a 19-year-old can navigate it without legal counsel and that the institution's response is constrained by something stronger than goodwill. It requires that the institution be willing to discontinue tools that fail equity audits, even when discontinuation imposes a sunk cost on a procurement decision someone is professionally invested in defending. None of these requirements is exotic. All of them are, at present, rare.

The French Inserm overview of AI's impact on education is useful here because it is unusually direct about the empirical question: the evidence that AI improves educational outcomes is real but conditional, and the conditions under which the improvements appear are mostly conditions the average institution is not providing [24]. The Illinois curriculum-and-training initiative is one of the more thoughtful institutional responses, in that it pairs the rollout of tools with the development of faculty capacity to think critically about them [4]. The instructional-design specialist programs that are now appearing across

[24] Quel est l'impact de l'IA sur l'éducation ?

[4] AI Curriculum and Training

the sector [6] are another signal that some institutions are taking the labor and expertise question seriously, though it remains to be seen whether the graduates of those programs will be hired in numbers sufficient to matter or treated, as instructional designers have historically been treated, as a service function rather than a strategic one.

[6] EdS in Instructional Design | AI Technology

The deeper point is the one Benjamin makes about poverty and educational attainment: it is not the facts that elude us. The evidence on what produces equitable educational outcomes is over a century old and unambiguous, and it has very little to do with which tool is licensed and a great deal to do with whether the institution is willing to distribute resources so that students can actually learn. AI does not change this calculus. It only changes the speed at which the institution can pretend it has.

The Choice the Institution Is About to Make

What is happening to higher education as AI arrives, then, is something more specific than "transformation" and more consequential than "disruption." The institution is being offered a deal: accept a class of tools whose benefits accrue to the already-advantaged and whose costs fall on the already-disadvantaged, in exchange for a politically convenient story about innovation and access. The deal is being offered by vendors with strong commercial incentives, accepted by administrators with strong reputational incentives, and absorbed by faculty and students who have not been meaningfully consulted. The accommodation trap is not a side effect of this deal. It is the deal's structure.

The exit from the trap is not a refusal of AI. The tools that work do real pedagogical good, and the institutions that figure out how to use them — Harvard's tailored physics tutor is one model, the human-AI collaborative configurations another — will offer their students something valuable. The exit is a refusal of the procurement frame: the assumption that adopting AI is fundamentally a question of which vendor to sign with rather than which pedagogical commitments to defend, which equity audits to fund, which redress mechanisms to build, which contingent labor to make permanent so that the tailoring work the successful cases require can actually be done at scale.

The institution that will make this transition is the one that has already begun to treat AI as a governance question rather than a procurement question, and that has located the governance authority somewhere other than the vendor's white paper. The institution that will not — and there will be many — will discover, slowly and then quickly, that the equity metrics it has been publishing are now being produced by systems it does not understand, on terms it did not set, in service of outcomes it would not, on reflection, endorse. The Delphi experts have already named the practices that distinguish the two paths [13]. The faculty members who have begun to redesign their courses around productive friction rather than algorithmic frictionlessness [27] have already begun to walk one of them. The question is whether the institution as such — the budgets, the contracts, the strategic plans, the accreditation responses — will follow them, or whether it will continue to mistake the procurement decision for the educational one and call the resulting gap an unfortunate edge case.

[13] Governing generative AI in higher education: a global Delphi ...

[27] The case for friction in AI-mediated information seeking and learning

The deaf graduate student whose captions render her advisor as a wall of nonsense is not an edge case. She is the test the institution is currently failing, and the test is not going to get easier.

References

1. A Competency Framework for Medical AI Education: Mixed Methods Study
2. Adelphi University accused a student of using AI to ... - Newsday
3. AI competency framework for teachers
4. AI Curriculum and Training
5. AI Governance in Higher Education: A course design exploring regulatory ...
6. EdS in Instructional Design | AI Technology
7. Faculty Ditch AI Bans: Study Shows Policy Shift - AcademicJ…
8. Falsas alarmas de vigilancia con IA han provocan castigos y arrestos ...
9. Frontiers | Student engagement with AI tools in learning: evidence from ...
10. GenAI in Higher Education, Legitimacy and Laziness
11. Generative AI in Higher Education
12. Generative AI Reduced Study Time on Math Problems and ...
13. Governing generative AI in higher education: a global Delphi ...
14. Impact de l'IA générative sur la « pensée critique »

15. Inteligencia Artificial y Pensamiento Crítico en Educación: Marcos ...

16. Intelligence artificielle : l'université peut-elle sanctionner sans règle

17. Karen Lumsden, PhD's Post

18. OpenAI s'étend dans l'enseignement supérieur indien via des ...

19. PDF AI and Graduate Teaching Labor: Reshaping Workload, Autonomy, and ...

20. PDF Perspective Étudiante Sur Les Systèmes D'Intellige Nce Artificielle ...

21. Preserving learning in the age of AI shortcuts — Harvard Gazette

22. Professor tailored AI tutor to physics course. Engagement doubled.

23. Public have more fear than hope on AI and future of work, ...

24. Quel est l'impact de l'IA sur l'éducation ?

25. Risk, Retention, and the Algorithmic Institution: Artificial Intelligence as a Policy Response to Higher Education in Crisis

26. School AI surveillance like Gaggle can lead to false alarms, arrests ...

27. The case for friction in AI-mediated information seeking and learning

28. The Impact of AI on Students' Reading, Critical Thinking, and Problem ...

29. The impact of an AI Digital Teacher on human-AI collaborative learning in higher education

30. The LOGOS Framework: A Five-Level Taxonomy of Human Cognitive Agency in AI-Assisted Assessment

31. Un étudiant sur 3 transgresse les règles à l'aide de l'IA