A FIELD NOTE FROM DIGICATION

Taking student AI *resistance* seriously.

Student resistance to AI is real, varied, and often arrived at through careful thinking. What follows is an inventory of the concerns these students raise, grouped by how addressable each one looks.

22 concerns · 2 cross-cutting observations · For students and faculty

TWO THINGS THAT FRAME EVERYTHING ELSE

Cross-cutting observations

Observation 01

Symmetry, or nothing.

When faculty generate course materials with AI and grade papers with AI rubric tools while telling students they cannot use AI, the message students receive is not pedagogical. It is a status assertion: I get to use this, you do not. The intellectual act students perform is then evaluated by a machine in seconds, which is an insult regardless of the resulting grade.

A student who has noticed the asymmetry stops giving the institution the benefit of the doubt on any other policy. The only stable position is symmetric: if student work must be human, faculty evaluation must be human. Institutions that split this have already lost the argument with their most thoughtful students.

Observation 02

Frontier exposure, or no credibility.

A faculty member or administrator who has only used the free version of a chatbot from 2023 is in roughly the position of a portrait painter being asked to make policy about photography in 1850 without having ever picked up a camera. They are entitled to their views; they are not in a position to set policy or to teach the medium.

What this requires is not that everyone become an AI enthusiast. It is that the people making decisions and modeling behavior for students have sustained, serious, agentic engagement with frontier AI under conditions of proper data privacy. Most institutional AI policy today is being set by people analogous to the painter who refuses to pick up the camera. Until that changes, much of the student frustration below is justified.

THE CATALOG

The 22 concerns

Learning & Cognition

№ 01

● WHERE SOLUTIONS EXIST

Skill atrophy

Research links frequent AI use to weaker critical thinking. Students worry about becoming fluent in prompting while losing the ability to articulate their own reasoning.

THE OBSERVATION

The research is real, and the trend it tracks is concerning. What the research is actually measuring, though, is one specific mode of AI use: students treating AI as an answer machine, accepting outputs without challenge, skipping the cognitive work the assignment was meant to develop. That mode does erode the thinking it claims to help.

A different mode of use is also possible. When AI is used as a questioner that challenges your reasoning rather than completing it, when it is asked to surface what you missed and stress-test what you concluded, the activity becomes metacognitive rather than substitutive. The research hasn't caught up with this distinction yet, but the pedagogical pattern is teachable and is used widely in institutions doing this work seriously.

The implication for students is that how you use AI matters more than whether. The implication for faculty is that assignments designed for AI-as-questioner can preserve, and sometimes strengthen, the developmental work AI-as-answerer would erode.

№ 02

● WHERE SOLUTIONS ARE EMERGING

Hollow learning experience

Students describe AI-assisted writing as feeling empty in a way pre-AI work did not. No grade improvement fills the absence.

THE OBSERVATION

The hollow feeling is real. It cannot be argued away. What can be said honestly is that the feeling tracks a specific mode of use: when AI is doing the hard work the assignment was meant to require, the experience hollows out, because the student is no longer doing the thing that gives the work its meaning.

When AI preserves the struggle, sharpens it, and is used to interrogate rather than to deliver, the felt quality changes. This is not something to argue with a student; it is something they have to experience.

The risk in claiming that AI can be used well is sounding like the solution is easy. It is not. It requires faculty to design differently, students to engage differently, and institutions to model the difference. The hollow feeling will keep showing up wherever those pieces have not come together.

№ 03

● WHERE SOLUTIONS EXIST

Loss of productive struggle

If AI does the hard part of an assignment, the student passes but doesn't develop. The skill the assignment was meant to build lives in the doing.

THE OBSERVATION

Students who feel this are protecting something real about their own education. The skill is in the struggle.

A useful comparison is what happened to financial education when the spreadsheet arrived. Before the spreadsheet, accounting and finance pedagogy at most levels stopped at arithmetic. The spreadsheet made it practical to teach what-if modeling, multi-variable scenario analysis, hypothesis-driven business reasoning. The struggle did not disappear; it relocated to harder, more useful questions. Teachers who taught for the post-spreadsheet world prepared students for the work that actually existed. Teachers who clung to hand-ledger pedagogy did not.

AI poses the same choice for many fields. The thinking has to relocate to where AI cannot complete it. That requires deliberate work from the institution. Without it, the floor falls out and nothing builds in its place. With it, the ceiling rises.

№ 04

● WHERE SOLUTIONS EXIST

Why am I paying for this

If AI can write the paper assigned in 2022 and faculty grade it with AI, tuition purchases a closed loop between two AI systems. The value question gets sharper every semester.

THE OBSERVATION

The question is not rhetorical and it deserves a real answer.

Institutions need to be able to articulate what tuition buys that AI alone cannot produce: mentorship from people who know the field, community with peers doing the same hard things, credentialing that reflects actual development, integrated human feedback at moments that matter, and a dramatically higher bar of what is expected because the old bar is now producible by software.

Institutions that do this work justify the cost. Institutions that do not, cannot. This is the work of the next five years in higher education, and it is the work that will separate institutions that thrive from institutions that do not.

№ 05

● WHERE SOLUTIONS EXIST

Metacognitive loss

Using AI removes the expertise needed to evaluate AI. Students become fluent in prompting but lose the capacity to judge what the tool produces.

THE OBSERVATION

Metacognition (the capacity to think about your own thinking) is not eroded by AI use; it is eroded by a particular mode of AI use. The same tool that bypasses your reasoning can also be asked to interrogate it.

Prompting AI to critique your argument, surface your blind spots, play devil's advocate, and stress-test your conclusions is a metacognitive exercise. It is hard to do alone and practical with a good tool.

This is core to how some institutions are teaching with AI now. It is also exactly the work reflection-based pedagogy was already doing, with AI as a thought partner rather than an answer machine.

Ethical

№ 06

● WHERE THE CONCERN STANDS

Environmental impact

Data centers consume water and energy. AI's growth is accelerating both. Students whose institutions have climate pledges see the contradiction.

THE OBSERVATION

The per-query numbers get weaponized in both directions, and the major AI companies have not released the kind of total-footprint disclosures that would settle the math. The opacity is itself a problem. The local impact, however, is unambiguous: data centers sited in water-stressed regions impose real costs on their neighbors.

WHERE ACTION LIVES

> Pushing for vendor transparency on energy and water use, because the absence of credible data is itself a problem we can name. Procurement pressure from institutional buyers, who can ask vendors to disclose before signing. Local civic engagement where data centers are sited in vulnerable regions, including supporting communities that are negotiating directly with operators. Student research projects on AI infrastructure decisions in your region.
>
> Nobody solves climate change by individual procurement choices. The aggregate trajectory is set by the largest decisions in the industry and by policy. Caring about the issue does not produce a clean win. It does produce visibility, and visibility is what eventually moves the policy and procurement levers that matter.

№ 07

● WHERE THE CONCERN STANDS

Labor exploitation

Workers labeling AI training data and moderating its outputs are often paid poorly and traumatized by the work. The pattern has now broadened from the Global South into US contract labor.

THE OBSERVATION

Workers doing the difficult and often traumatizing labor of labeling data and moderating content sit underneath the AI products students are being asked to use. The most extreme conditions persist in parts of Africa and Southeast Asia, where wages are lowest and protections weakest. More recently, similar exploitation patterns have surfaced in the United States, where contractors hired to do AI knowledge work are paid far below the value being extracted from their labor. The geography varies; the pattern does not.

WHERE ACTION LIVES

> The remedies for this kind of harm are long-arc remedies, and they have a poor record of fast wins. Labor regulation is generational work. Disclosure requirements get passed and then weakened. Procurement standards at one university take years to write and may not survive a change in administration. Investigative journalism makes a practice visible and then competes for attention against everything else.
>
> What that looks like for someone who cares: pick a piece of the system and contribute to it. Student researchers documenting supply chains have fed earlier waves of reform; that work matters whether or not any single project leads to policy change. Procurement standards can be pushed at one's own institution, even when the result is partial and the institution drags its feet. Voting matters at the slow tempo at which labor law actually moves. Faculty can support this directly, including by using AI as a research tool to help students investigate the supply chains underneath the tools we all use. None of this guarantees change. The alternative, disengagement, guarantees that nothing changes from the side of the people who care.

№ 08

● WHERE THE CONCERN STANDS

Copyright and consent

AI companies absorbed creative work without consent to train their models. Unlike search, which pointed users back to sources, AI synthesis returns nothing to the creator.

THE OBSERVATION

The structural difference from search engines is meaningful. Search drove traffic back to publishers and creators. AI synthesis absorbs the work into model weights and returns nothing. The legal question is unsettled and lawsuits are ongoing; the ethical question depends on whether you read training as learning-like or copying-like, and reasonable people disagree.

The historical parallel with Google in 2000, when publishers also did not consent and the value exchange was established later, suggests humility about predicting how this resolves. It also suggests that the resolution will not happen on its own.

WHERE ACTION LIVES

> Following and supporting the lawsuits and policy debates already in motion, because their outcomes will set the rules. Supporting opt-out mechanisms for creators and the platforms that respect them. Pushing for compensation models in the institutional licensing of AI tools. Faculty using current cases as teaching material rather than treating the question as settled. Voting for representatives who back creator-side reforms.
>
> The work is slow and many of the early outcomes will be partial. The conversation that produces eventual norms is happening now and is shaped by who participates.

№ 09

● WHERE THE CONCERN STANDS

Bias and representation

Models reflect and amplify biases in their training data, skewed toward Western, English-language content. Students from marginalized groups see this most clearly.

THE OBSERVATION

The tool is a mirror of society's biases, sometimes magnified by training and deployment choices. There is no clean technical fix. Students from marginalized groups are not wrong to point at it.

WHERE ACTION LIVES

> Becoming skilled at counteracting bias in practice, including asking for multiple perspectives, questioning default framings, and comparing outputs from different models. Documenting biased outputs in high-stakes contexts so the patterns are visible. Pressing institutions to evaluate AI tools for bias before deploying them in grading, admissions, or advising decisions. Supporting models trained on more diverse data and in more languages, including the academic and open-source efforts working on this. Research on bias in AI used in education itself, conducted by the students and faculty affected.
>
> None of this fully solves the underlying problem. All of it shifts the practice toward something better than the passive acceptance of biased outputs in places where they cause real harm.

№ 10

● WHERE SOLUTIONS ARE EMERGING

Privacy and data extraction

Every prompt is potentially training data. Free tools rarely promise otherwise, and most students default to those.

THE OBSERVATION

The concern is structural, not paranoid. For informed users with access to enterprise-grade tools, this is largely solvable. Major providers offer enterprise terms that do not train on user data, and the biggest players face serious legal liability for violating those commitments. Educational technology can be built on those terms by default.

The harder piece is the equity gap underneath. If you have money, your data is safe; if you do not, it is not. That is a structural concern in its own right. The answer is not to tell individual students to upgrade. It is for institutions and the educational tools they choose to deliver enterprise-grade protection by default, so the student does not have to know to ask. This is a concrete design choice. It is the kind of choice educational software companies should be making, including this one.

№ 11

● WHERE THE CONCERN STANDS

Deepfakes and targeted harm

AI has enabled a scale of targeted harm, especially against women and girls, that did not previously exist. Nudify apps, voice cloning, manufactured intimate imagery.

THE OBSERVATION

The harm is real, the victims are disproportionately female and disproportionately young, and the response from the broader AI ecosystem has been slow and uneven.

WHERE ACTION LIVES

> Supporting legislation requiring AI labeling and prohibiting non-consensual intimate imagery; some states have passed laws in the last two years and more are pending. Reporting harms when they occur, supporting victims, and treating the issue as serious rather than minimizing it as a side effect of progress. Digital literacy education that takes the threat seriously rather than treating deepfakes as a curiosity. Pressure on platforms for better moderation and faster takedown of non-consensual content. Voting for representatives who treat AI-enabled harassment as the urgent issue it is.
>
> The frustrating truth is that the technical capabilities will keep advancing faster than the legal and platform response. The slow work of building legal frameworks and cultural norms is the answer anyway, because the alternative is to leave victims with no recourse.

№ 12

● WHERE SOLUTIONS EXIST

Academic integrity as personal value

Many students feel they're cheating when they use AI, even when permitted. They watch classmates cheat and get equal or better grades. The unfairness corrodes their willingness to do honest work.

THE OBSERVATION

This is a character concern, not a rule-compliance concern, and it deserves to be taken on its own terms.

The answer is to redesign the work so the cheating question dissolves. When an assignment can be adequately produced by AI alone, the assignment has a problem, not the student. When AI-assisted work requires judgment, integration, iteration, and genuine thinking that AI alone cannot produce, the integrity question reframes: everyone uses AI, and everyone also thinks, and the thinking is what is being evaluated.

This is not a cosmetic change. It requires faculty to rebuild assignments around cognitive tasks that AI alone cannot complete. The institutions that do this work solve the integrity problem at its root.

Coercion

№ 13

● WHERE SOLUTIONS ARE EMERGING

Institutional partnerships students didn't consent to

Universities sign deals with AI companies without consulting faculty or students. The partnership is done to them, not with them.

THE OBSERVATION

This is a governance failure more than a technology failure, and the distinction matters because governance failures have known fixes. Real consultation processes. Transparent decision-making. Faculty voice in tool selection. Student representation in the conversations that lead to institutional AI commitments.

The opportunity is open. Higher education has not yet built the norms it needs around AI procurement and partnership. The institutions that do this well over the next few years will be the ones that consult before signing, disclose what they have agreed to, and explain why.

№ 14

● WHERE SOLUTIONS EXIST

Incoherent messaging

'AI will reshape your career, but don't use any AI' is contradictory. Policies vary by class, by semester, by professor.

THE OBSERVATION

Coherent AI policy is not a hard problem in theory. It is a small task force, a clear institutional stance, a communication plan, a set of operating principles that hold across courses.

The reason most institutions have not done it is not that the problem is intellectually difficult. It is that the work is politically uncomfortable, requires faculty consensus, and forces institutions to take positions that some constituency will not like.

The work is solvable. It is just being avoided. The institutions that do the work will get an immediate improvement in how their students experience their AI policy, simply by making it consistent.

№ 15

● WHERE SOLUTIONS ARE EMERGING

The get-left-behind pressure

Telling students they must use AI or be unemployable is a form of coercion, especially for students with ethical objections.

THE OBSERVATION

The pressure is real and it's reaching the job market: entry-level roles in AI-exposed fields are visibly contracting.

What is durable is not prompting (which is becoming commodity within two years) but the capacity to design AI-human-agent systems, to reason about when to involve which actor, and to hold quality standards across mixed workflows. That capacity is teachable, and it can be taught to students who are not uncritical AI enthusiasts. It does not require ideological commitment. It requires sustained practice with the tools in a context where the practice has stakes.

Students who develop these meta-skills enter a labor market that still needs them, even where prompt-only roles are disappearing. Institutions that teach this well are doing the most concrete thing they can do for student employability.

№ 16

● WHERE SOLUTIONS EXIST

False-positive AI detection

Students with clean writing styles get flagged by unreliable AI detectors. Honest students feel presumed guilty by surveillance that does not work.

THE OBSERVATION

AI detection tools are genuinely unreliable, and using them as a disciplinary instrument is unsafe.

More importantly, the existence of AI detection as a strategy is a pedagogical admission of failure. If a detector matters, the assignment could be produced by AI. The fix is assignment design, not surveillance. Many institutions are arriving at this conclusion and phasing out detectors. The trend is moving in the student's favor, and it should.

№ 17

● WHERE SOLUTIONS EXIST

Inconsistent policies across classes

Every semester, every class, new rules. Students navigate compliance puzzles on top of the actual learning.

THE OBSERVATION

The cost is borne by students; the dysfunction is institutional.

This is solvable through real institutional coordination, which is the kind of easy that is not actually easy. It requires a faculty body to come to a shared position, communicate it, and hold to it. It requires administrators to back the position when individual faculty want to depart from it. It requires the political work that most institutions defer.

What students should expect is consistency. What most institutions deliver is fragmentation. The gap between expectation and reality is what produces the frustration. The fix is procedural and known.

Identity

№ 18

● WHERE SOLUTIONS ARE EMERGING

Protection of voice

For students who write, voice is the asset being developed. Outsourcing writing to AI threatens the developmental process, even when outputs look fine.

THE OBSERVATION

Writing is not only production; it is often the process by which thought gets refined. You do not always know what you think until you have written it.

The pre-industrial clothing analogy is useful here. Hand-making clothes once taught physics, materials, and proportion. Industrial production made hand-making economically optional but did not erase its developmental value, which is why some curricula preserved it as pedagogy long after the economic case had collapsed. AI poses the same choice for writing. The question is not whether AI can produce prose. It is whether the developmental practice of writing through one's own thinking remains worth teaching, and on what terms.

The answer is yes, but with curricular commitment that has not yet been broadly made. Writing-as-thinking has to be defended as a pedagogical practice, not as a production technology. That defense is being built; it is not yet finished.

№ 19

● WHERE THE CONCERN STANDS

Corporate capture of the tools of thinking

A handful of companies shape the tools that shape how a generation of people think. Open-source alternatives exist but are dwarfed by the commercial frontier.

THE OBSERVATION

The concern is shared by tech ethicists, antitrust regulators, and democratic theorists. It is not fringe. The dominant tools are owned by corporations whose incentives are not the same as the public interest.

WHERE ACTION LIVES

> Supporting open-source and academic AI efforts as users and as funders. Pushing universities to collaborate on public-interest LLMs with transparent training data and democratic governance; universities have the intellectual capacity, and the obstacle is coordination and institutional courage, neither of which is technical. Advocating for antitrust action and platform regulation. Voting for representatives who treat tech concentration as the structural problem it is. Choosing alternatives where they exist and are good enough.
>
> None of this dissolves the underlying concentration. The concentration is real and durable. What action does is preserve the conditions under which non-corporate alternatives can develop, and create pressure for the eventual antitrust and regulatory action that genuine alternatives would require.

№ 20

● WHERE SOLUTIONS ARE EMERGING

Social signaling from faculty

When professors signal 'AI will rot your brain,' students absorb that. Uncritical enthusiasm creates the opposite problem.

THE OBSERVATION

Both stances fail students who are trying to figure out what is true.

The useful stance is specific and evaluative: I use AI for these things, not for those, because of these reasons. Students can absorb that and develop their own version of it. They cannot do much with 'AI is bad' or 'AI is great' except pick a side.

Faculty modeling matters in both directions. The pedagogical work is to develop a specific, defensible personal AI practice that the faculty member can describe and defend. The institutions that train toward this stance, rather than letting individual faculty improvise both directions, are giving students the model they need.

№ 21

● WHERE SOLUTIONS EXIST

Job market anxiety, inverted

Why enthusiastically adopt the technology that is erasing the career path you prepared for?

THE OBSERVATION

The grief is legitimate. So is the frustration that the technology arrived without anyone's consent.

What survives the contraction is more, not less, dependent on the durable skills education has always claimed to teach: critical thinking, working with others, communication, growth mindset, the capacity to do hard things and recover from failure. These matter more in an AI-saturated workplace, not less.

The problem is how students demonstrate them. An eight-page paper submitted to a learning management system is read by no one, including the AI agents that increasingly screen candidates. The paper does not show critical thinking; it produces an artifact that someone hopes contains critical thinking. The demonstration is broken.

What can work: structured reflection, and the ability to share it. Portfolios that show what a student brings beyond what prompting can produce. Evidence of thinking, growth, and judgment, organized so that a hiring manager or their AI screener can actually see it. This is workable today and is among the most concrete answers higher education has to the job-market question.

Military & State

№ 22

● WHERE THE CONCERN STANDS

Military, surveillance, and authoritarian use

The same generative AI being marketed in education is being deployed by intelligence agencies for targeting, by immigration enforcement to surveil civilians, and by authoritarian regimes to suppress dissent.

THE OBSERVATION

The companies selling AI for education are, in many cases, the same companies sitting on those military and surveillance contracts. The line keeps shifting toward the latter, and students who are paying attention notice. The AI industry's relationship with the security state is structural, not incidental.

WHERE ACTION LIVES

> Supporting export controls and disclosure requirements on AI used in lethal or surveillance applications. Investigative journalism, including student journalism, on dual-use deployments of AI by US agencies and contractors. Choosing where to work after graduation, individually or collectively; employee organizing against military contracts has affected company behavior in recent years. Universities applying procurement standards to defense-adjacent AI use, refusing to do business with vendors whose other revenue streams the institution cannot defend.
>
> The action paths exist; they are slow, and they have a mixed record. The students who pursue them are doing the work that produces eventual norms, the same way previous generations did this work around chemical weapons, nuclear proliferation, and surveillance technology. The outcomes are not guaranteed. The work matters anyway.

THE BREAKDOWN

How the 22 concerns distribute

- **9 Where solutions exist:** Areas where pedagogy, tools, or institutional posture provides a real answer. May still require execution work, but the path is known.
- **6 Where solutions are emerging:** Areas where the direction is real and the work is underway, but the finished answer doesn't exist yet. Workable, not solved.
- **7 Where the concern stands:** Areas where the concern is genuine, hard, and not addressable at the scale of any one company. Action paths exist; easy wins do not.