




Higher Ed AI Policy Compass

A framework for understanding where your institution stands on AI governance and what that stance is actually producing.



balloon is inflated, it starts to begin to oxidize and will break" (Violet 2010). Biodegradable balloons are made of natural latex and...




Table of contents

- 03** Section 1: About this resource
- 04** Section 2: Three postures. One spectrum.
- 05** Section 3: The three postures
- 09** Section 4: Pressure-testing your current posture
- 11** Section 5: The implementation bridge
- 14** Section 6: Using the compass over time





Section 1

About this resource

Higher education institutions make decisions about AI every day—in policy committees, department meetings, and faculty conversations—well before any formal review cycle. Some of those decisions are deliberate. Many are not.

This resource is designed to support the deliberate ones.

The Higher Ed AI Policy Compass is a framework for understanding where your institution currently stands on AI governance and what that stance is actually producing. It helps institutions move from reactive policy-making to intentional governance. By mapping three common policy postures, it surfaces the tensions each one generates and offers a set of reflection questions to help you pressure-test where you are before deciding where to go.

What this is not

This is not a scoring tool. It will not tell you which posture is right. It won't produce a recommendation or grade your current policy. What it will do is give you language for what you are already seeing, along with a structure for conversations that might otherwise remain fragmented.


Think of it as a compass, not a map. It helps you orient, but it doesn't chart the route.

A compass doesn't tell you where to go.
It helps you to understand where you are.

Who this is for

This resource is designed for academic leaders and governance stakeholders: provosts, CIOs, deans, and the cross-functional committees responsible for AI policy decisions. It assumes you're already in motion, that some version of an AI policy exists or is under development, and that you're looking for a way to assess it with more clarity and confidence.





Section 2

Three postures. One spectrum.

Before you can assess your policy, it helps to understand the terrain.

Most institutional AI policies, regardless of how they're written or framed, reflect one of three underlying postures: **Restrictive**, **Hybrid**, or **Permissive**. These are not rigid categories so much as orientations. They describe the assumptions your policy makes about risk, trust, and the role of AI in the learning process.

At a high level:




A posture is the set of conditions your policy creates for faculty, students, and administrators as they navigate AI use in real time. Two institutions can use nearly identical language and land in very different places, depending on how that language is interpreted, enforced, and communicated.

Understanding your posture means understanding more than what your policy says. It means understanding what your policy *does*.

What makes a posture

Each posture reflects a different answer to the same underlying question: What do we believe about AI, and what does that belief require of students, faculty, and administrators? Posture shapes everything downstream, from how policy is interpreted to how trust is communicated, how tensions surface, and how governance adapts over time.





Section 3

The three postures



Restrictive

The core logic

Restrictive postures begin from a protective premise in which AI use risks compromising the integrity of learning and assessment. The logic then becomes protective. Limit AI's role to preserve the validity of assessments and the authenticity of student work. This posture typically appears as clear prohibition language, with defined consequences for violations and an emphasis on detection and enforcement.


What it looks like in practice

Faculty experience clarity, at least initially. There is a defined line, and most welcome having it. Students operate under a presumption of compliance, though some seek workarounds. Administrators spend meaningful energy on enforcement: detection tools, academic integrity processes, and policy interpretation in edge cases. The conversation shifts from learning to compliance.

Tensions that emerge

The tensions in a restrictive posture are predictable. Detection tools are imperfect and frequently contested. The burden of proof often falls on students in ways that can feel arbitrary and erosive of trust. Faculty are asked to serve as compliance officers in addition to educators, which can feel like a role most did not sign up for. There is also a subtler tension: Restrictive postures can inadvertently suppress legitimate uses of AI, including uses that could support students who struggle with writing mechanics, language barriers, or accessibility needs.





Restrictive (Continued)

The integrity signal

In a restrictive posture, trust is demonstrated through absence. The implicit message to students is: Prove you didn't use AI. Academic integrity becomes synonymous with avoidance, which is a narrow and increasingly difficult standard to enforce.

When trust is demonstrated through absence, integrity becomes a surveillance problem instead of a learning one.



Hybrid

The core logic

Hybrid postures operate on the belief that context matters. What is appropriate in a capstone research seminar is not appropriate in a first-year writing course. Rather than establishing a single institutional stance, hybrid policies delegate authority (to departments, individual faculty, or specific assignment types) to determine where and how AI can be used.

The logic is nuanced: Learning needs vary, and a one-size-fits-all policy serves none of them well.


What it looks like in practice

Flexibility is the defining feature. It is also the defining challenge. Faculty carry significant interpretive and decision-making weight. Students navigate different rules across different courses, sometimes within the same semester. Administrators manage a policy that's less a rulebook than a framework, which means interpretation questions surface constantly.

In practice, hybrid postures require institutional infrastructure that many institutions do not yet have: shared language, regular communication, and sufficient trust among stakeholders to enable local decision-making.

Hybrid governance succeeds or fails on coordination capacity. Without shared norms and communication infrastructure, flexibility becomes fragmentation.





Hybrid

(Continued)

Tensions that emerge

Inconsistency is the most common friction point. Students experience the policy differently depending on the instructor, the department, or the assignment. Faculty face a version of decision fatigue: the burden of making the right call, without always having the support or clarity to do so confidently.

The middle ground can also become no ground at all. Without a strong coordination infrastructure, a hybrid posture risks becoming a de facto permissive one. This happens not because the institution chose it, but because enforcement became too complicated to sustain.

The integrity signal

In a hybrid posture, trust is demonstrated through transparency. Students are asked to disclose their AI use and explain it. Integrity becomes a conversation about process transparency rather than the output, which is more sophisticated and demanding.

Transparency is a higher bar than avoidance. It asks students not just to follow rules, but to understand their own reasoning.




Permissive

The core logic

Permissive postures start from a different premise: AI is a tool, and like other tools, learning to use it well is itself a valuable skill. Rather than restrict or contextually limit AI use, permissive policies focus on outcomes: what students produce, demonstrate, and understand.

The logic is pragmatic and forward-looking. The question shifts from “Did you use AI?” to “What can you do with AI?”





Permissive (Continued)

What it looks like in practice

Faculty experience significant creative latitude alongside significant interpretive responsibility. The policy asks them to design assessments that hold up even when students have access to AI. Students generally appreciate the freedom, but some find the lack of guardrails disorienting, particularly if they're still developing a sense of their own capabilities.

Administrators face a different challenge: What does consistency look like when the policy is intentionally hands-off?

Tensions that emerge

Equal access is the tension that surfaces earliest. The risk is not only unequal access to tools, but also unequal fluency in using them critically and strategically. Not all students enter with the same access to AI tools, the same fluency in them, or the same understanding of how to use them critically. A permissive posture can inadvertently reward students who are already advantaged while leaving others behind.


Assessment validity is the other pressure point. When AI can produce a serviceable draft in seconds, what are we actually measuring? Faculty are asked to answer that question course by course, assignment by assignment, which is appropriate but not easy.

The integrity signal

In a permissive posture, trust is demonstrated through engagement. Students are asked to show their thinking, not just their output. Integrity becomes visible in process: the iterations, the reasoning, the choices made with and around AI.

Permissive doesn't mean anything goes. It means the expectation has shifted from compliance to comprehension.





Section 4

Pressure-testing your current posture

These questions are meant to surface what you already know, so you can bring it into the room when it matters.

Diagnostic Reflection

By now, you likely recognize your institution, or something close to it, in one of these postures. Maybe you see a clean fit. Maybe you see a hybrid of postures operating simultaneously in different parts of your institution. Either is normal.

What follows is not a quiz. There is no score. These are questions to sit with, individually or with your governance committee, before you decide what, if anything, needs to change.

Alignment

Does your stated posture match how it's actually playing out?


Policy and practice rarely align perfectly. What is written in a policy document reflects the institutional intent. What happens in classrooms, on assignments, and in academic integrity proceedings reflects reality. Where do you see the biggest gaps? Are different parts of your institution operating under different de facto postures? If an external observer watched your classrooms and integrity cases for one semester, what posture would they conclude you hold?

Trust

What does your current posture signal to students? To faculty?

Every policy communicates something beyond its literal content. Restrictive policies can signal vigilance and, at times, suspicion. Permissive policies can signal confidence and occasionally indifference. What does your policy say to the people most affected by it? Is that the message you intend?





○ **Lived experience**

What are faculty and students actually experiencing?

This question is harder to answer than it sounds. Formal feedback mechanisms often capture the voices of the most satisfied and the most frustrated. What do you know about how your policy is landing with the people navigating it daily? What would you need to learn to understand that more fully?

○ **Stability over time**


Is your current posture durable, or is it already showing cracks?

AI capabilities are evolving faster than most policy cycles can accommodate. A posture that felt adequate a year ago may be straining under new tools, new student behaviors, or new faculty concerns. How confident are you that your current approach can hold up through the next policy review cycle?

A note on self-assessment

Honest self-assessment in higher ed governance is harder than it looks. Institutional inertia, political dynamics, and the pressure to appear decisive can all distort perception. If you are using this resource with a committee, consider establishing shared norms before you begin, including permission to name what isn't working.





Section 5

The implementation bridge

From posture to practice

Choosing a posture is just the beginning of the journey. Even a well-designed policy, clearly communicated and broadly accepted, does not automatically produce consistent implementation. Something happens between the policy document and the lived experience of students and faculty, and understanding that gap is where governance becomes real.

What follows are four patterns that appear across postures. Instead of thinking about them as problems to be solved, it can be more productive to view them as structural features of how policy operates in complex institutions. Naming them helps.

Posture sets conditions. It doesn't determine outcomes.



Pattern 1

Policy language → lived interpretation

What is written in a policy document is not what most people read. Faculty interpret policy through the lens of their discipline, their past experience, and their intuitions about what the institution actually values. Students interpret it through what they hear from peers, what their instructors emphasize, and what they observe being enforced. The distance between policy language and lived interpretation is not a communication failure. It's a feature of how complex organizations work. The question then becomes how to manage it deliberately.




Pattern 2

Intent → signal

What you intend to communicate through your policy is rarely what people hear. A restrictive policy designed to protect academic integrity can lead to institutional distrust of students. A permissive policy designed to honor faculty autonomy can be perceived as administrative indifference to learning quality.





Pattern 2 (Continued)

The signal your policy sends is shaped as much by context, history, and relationships as by the words on the page.

Policies do not operate in a vacuum. They amplify existing levels of trust, or mistrust, already present in the institution.

Understanding the gap between intent and signal requires listening, which is a governance practice and not a communication strategy.



Pattern 3

Autonomy → coordination load

More flexibility in a policy creates more decisions that need to be made locally. More locally made decisions create more inconsistency across courses, departments, and student experiences. More inconsistency increases the burden on students, who must navigate different expectations with no clear map.

This is an argument for building the coordination infrastructure that flexibility requires. Without it, autonomy becomes a tax on the people least equipped to pay it.

In practice, that tax is often paid by students navigating inconsistent expectations across courses.



Pattern 4


Governance → renewal confidence

Institutions that build structured review cycles into their governance model are better positioned to adapt when things change.

Those that respond to AI policy primarily when a high-profile incident forces a reaction are more likely to overcorrect and to undermine the trust they're trying to protect.

Renewal confidence, the capacity to revisit and revise without crisis, is a governance asset, and it doesn't happen by accident.





Cross-posture insight: Misalignment is predictable

Certain tensions appear across every posture:



**Detection
challenges**



**Inconsistent
faculty
enforcement**



**Student
confusion at the
edges of policy**

These aren't signs that your posture is wrong, but they indicate that you are operating in a genuinely complex environment. Naming predictable tensions in advance does not prevent them, but it does mean you're not surprised when they arrive. Institutions that aren't surprised are better positioned to respond thoughtfully than those that are.


Closing note

There is no frictionless AI policy. The goal is not to eliminate tension, but to govern it intentionally.

Where alignment breaks down

The most common misalignments we observe are not between policy and intent. Rather, they're between intent and signal. Institutions often underestimate how much weight the communication and culture surrounding a policy carry. The policy document is one signal. Faculty behavior, institutional history, and how enforcement decisions are handled are others. The most important thing to remember is that *all of them* matter.





Section 6

Using the compass over time

Governance is ongoing

The compass is not a one-time diagnostic. It's a reference point that is useful at the start of a policy review, useful in the middle of a governance conversation, and useful when something in your environment shifts and you need to reorient quickly.

When the compass is most useful


There are moments when this resource is particularly well suited to the work at hand:

- ☑ When you're entering a formal policy review cycle and you want to begin from a clear-eyed assessment of where you actually are, not where you think you are.
- ☑ When a high-profile incident or external pressure is pushing your institution toward a reactive posture shift and you want to ensure any changes are intentional.
- ☑ When new stakeholders join a governance committee and need context on where the institution has been and what it's navigating.
- ☑ When you sense that your stated posture and your lived practice have drifted and you want to name that gap before it becomes harder to address.


How to use this resource internally

This document works best as a conversation tool, not a read-alone artifact. Consider distributing it in advance of a governance meeting and using the diagnostic questions in section 4 as discussion prompts. Different stakeholders will recognize your institution differently, and that variation is often where the most useful insight resides.





HOW TO USE THIS IN A 60-MINUTE MEETING



A note on change over time

Your posture today may not be the same as your posture in two years. That's not failure. AI capabilities are evolving, student and faculty expectations are shifting, and the broader higher education landscape is still finding its footing. Intentional posture evolution driven by reflection and evidence (not crisis) is the mark of a governance structure that's working.

The institutions best positioned for that kind of evolution are the ones that have built the habits of honest self-assessment: asking hard questions, making space for disagreement, and staying curious about the distance between what they intend and what they create.

Closing note

The compass doesn't tell you where to go. It helps you know where you are so that wherever you go next, you go on purpose.

Intentional posture evolution driven by reflection and evidence (not crisis) is the mark of a governance structure that's working.

ABOUT THIS RESOURCE

The Higher Ed AI Policy Compass was developed by Superhuman to support institutional leaders as they navigate AI governance decisions. It is designed to be shared across governance stakeholders and adapted to your institution's specific context.





SUPERHUMAN

Superhuman for Education extends Grammarly's trusted writing support into a human-centered AI platform designed for teaching, learning, and working. It helps students write and learn with AI through visible, defensible processes, and provides faculty and staff with context-aware, agentic assessments and assistance embedded in their workflows, enabling transparent, controlled, agentic AI use that helps institutions innovate and meet today's challenges head-on.

[Learn more](#)