

University AI Policy Guidelines

A Comparative Research Brief for Higher Education Leaders

Extracted Policy Guidelines from Leading Global Universities

Prepared February 2026

Universities Covered:

Arizona State University | Stanford University | Harvard University | Cornell University | University of
Pennsylvania (Wharton) | MIT | Duke University | University of Sydney

Introduction

This document compiles and summarizes the official AI usage policies from eight leading universities worldwide. These policies are extracted directly from institutional guidelines published between 2023 and early 2026. The purpose is to provide a comparative reference for university leaders developing their own institutional AI frameworks.

Important note: University AI policies are evolving rapidly. Most institutions update their guidelines multiple times per year. The policies summarized here reflect the most current publicly available versions as of February 2026. Always consult the original source links for the latest versions.

1. Arizona State University (ASU)

Source: ai.asu.edu/digital-trust-guidelines and ai.asu.edu/policy-and-resources

1.1 Governing Framework

ASU operates under a "Principled Innovation" framework. Guidelines were reviewed and approved by Enterprise Technology Digital Trust, Cybersecurity, AI Acceleration, Learning Experience, and the Office of General Counsel.

1.2 Key Policy Provisions

Data Privacy and Security

- Do not share information that is internal, sensitive, or highly sensitive into generative AI tools. This includes student records, employee data, research data subject to agreements, financial information, and proprietary institutional data.
- Before using any AI tool, understand whether the tool can keep or use the information you submit and whether that conflicts with your other responsibilities.
- Only use AI tools that have completed the ASU VITRA (Vendor IT Risk Assessment) process and are approved for departmental use.

Academic Integrity

- The unauthorized use of ChatGPT and other generative AI technologies violates the University's academic integrity policy.
- Academic units and faculty determine whether student use of generative AI in their courses is permitted or prohibited. This must be stated in syllabi, announcements, and assignment instructions.
- If students use generative AI, they must cite it according to applicable citation guidelines (APA, etc.).
- Suspected use of generative AI in coursework is NOT sufficient evidence to begin a formal Academic Integrity investigation.
- The university does not recommend specific AI detection tools due to their unreliability.

Intellectual Property

- By submitting information to a generative AI tool, you may give up valuable rights to control who can use that intellectual property.
- Materials prepared using a generative AI tool may not be eligible for intellectual property rights if there is insufficient direct human involvement.
- Do not submit content owned by the Arizona Board of Regents, another institution, or a co-author without authorization.

Research Use

- Discuss the appropriateness of AI tools with co-investigators, collaborators, and field experts before starting any research project.
- Do not place federal, state, or ASU data into externally sourced generative AI tools.
- Do not place any US export-controlled data into any generative AI tools.
- Review award terms and conditions for AI-related standards before using AI in federally funded research.

Staff/Marketing Use

- AI tools are partners in work, speeding up concept generation and initial drafts, but they do not replace the expertise that ASU staff bring to the final product.
- Humans must iterate on AI outputs to ensure alignment with university branding and messaging.
- Be aware of bias AI might generate. AI users hold the responsibility to monitor and reduce input and output bias.
- Do not use AI to change people's actual physical appearance in images, create fabrications, or generate content in the style of a specific artist.

1.3 Unique Approach

ASU provides universal free access to ChatGPT Edu (with GPT-5) for every student, faculty, and staff member, with enterprise-level privacy protections ensuring data is not used to train AI models. The university has also launched an AI Playground, self-paced AI courses, and an AI Innovation Challenge that has activated over 500 projects.

2. Stanford University

Source: communitystandards.stanford.edu/generative-ai-policy-guidance

2.1 Core Policy (Board on Conduct Affairs)

- **Default rule:** Absent a clear statement from a course instructor, use of or consultation with generative AI shall be treated analogously to assistance from another person.
- **Prohibited use:** Using generative AI tools to substantially complete an assignment or exam (e.g., by entering exam or assignment questions) is not permitted unless the instructor explicitly allows it.
- **Disclosure:** Students should acknowledge the use of generative AI (other than incidental use) and default to disclosing such assistance when in doubt.
- **Faculty autonomy:** Individual course instructors are free to set their own policies regulating the use of generative AI tools in their courses, including allowing or disallowing some or all uses. These must be communicated in course syllabi.

2.2 Data Protection

- Students and staff should not provide any confidential or legally privileged information of Stanford or a third party to generative AI tools.
- Stanford provides an AI Playground with enterprise-level protections where files uploaded are not shared externally or used to train models.

2.3 Detection Tools

Stanford's Academic Integrity Working Group identified several issues with AI detection tools, including bias and occurrences of false positives and negatives. They found these tools unsuitable for high-stakes situations, especially as evidence in academic misconduct cases.

2.4 Assessment Recommendations

- For instructors who wish to limit AI use, the group recommends in-person formats such as oral exams and in-class writing assignments.
- Departments and multi-section courses are encouraged to establish consistent and enforceable AI policies that can be applied fairly across sections.

2.5 Staff/Communications Guidelines

- AI should be used to augment, not replace, human work.
- All employees communicating on behalf of Stanford are encouraged to engage in regular discussion of their AI practices with colleagues and supervisors.

- Regularly review and assess the models you are using to ensure they meet expectations and guidelines.
- You are responsible for ensuring that any generative model provides a license for outputs.

3. Harvard University

Sources: provost.harvard.edu/guidelines-using-chatgpt; oue.fas.harvard.edu/huit.harvard.edu/ai/guidelines

3.1 University-Wide Guidelines (Provost)

- **Data protection:** Do not enter data classified as confidential (Level 2 and above), including non-public research data, into publicly-available generative AI tools.
- **Responsibility:** You are responsible for any content that you produce or publish that includes AI-generated material. AI-generated content can be inaccurate, misleading, or entirely fabricated.
- **School policies:** Schools will develop and update their own policies. Faculty should be clear with students about permitted uses of generative AI.
- **Security:** All vendor generative AI tools must be assessed for risk by Harvard's Information Security and Data Privacy office prior to use.

3.2 Faculty of Arts and Sciences (FAS) Framework

FAS does not impose a single AI policy. Instead, it provides three template syllabus statements for faculty to choose from:

- **Maximally restrictive:** "We specifically forbid the use of ChatGPT or any other generative artificial intelligence (AI) tools at all stages of the work process, including preliminary ones. Violations of this policy will be considered academic misconduct."
- **Fully encouraging:** "This course encourages students to explore the use of generative artificial intelligence (GAI) tools such as ChatGPT for all assignments and assessments. Any such use must be appropriately acknowledged and cited."
- **Mixed approach:** "Certain assignments in this course will permit or even encourage the use of generative artificial intelligence (GAI) tools. The default is that such use is disallowed unless otherwise stated."

All three templates include: "It is each student's responsibility to assess the validity and applicability of any GAI output that is submitted; you bear the final responsibility."

3.3 Harvard Graduate School of Education

HGSE provides more specific guidance on permissible vs. impermissible uses:

- **Permissible:** Seeking clarification on concepts, brainstorming ideas, generating scenarios, AI-powered web search, having conversations with tools like ChatGPT to explore ideas, refine thinking, identify examples. Also acceptable for drafting emails not submitted as coursework.
- **Required:** For any permitted use, students must acknowledge and document that use, explaining what tools were used, prompts provided, and how output was integrated.

- **Faculty override:** Individual instructors may create course-specific policies that differ from and supersede these guidelines.

3.4 Assessment Innovation

Starting Fall 2025, Harvard deployed Respondus, a browser lockdown tool, for in-person exams to ensure students do not use AI unless the course asks them to. Faculty are discouraged from using AI detection tools due to unreliability.

4. Cornell University

Sources: it.cornell.edu/ai/ai-guidelines; teaching.cornell.edu/generative-artificial-intelligence

4.1 IT Guidelines

- **Confidentiality:** If you are using public generative AI tools, you cannot enter any Cornell information, or another person's information, that is confidential, proprietary, subject to federal or state regulations, or otherwise considered sensitive or restricted.
- **Public data:** Any information you provide to public generative AI tools is considered public and may be stored and used by anyone else.
- **Verification:** Always verify the information for errors and biases and exercise caution to avoid copyright infringement. Generative AI cannot understand what it produces; results are sometimes misleading, outdated, or false.
- **Privacy principles:** Notice, Choice, Accountability for Onward Transfer, Security, Data Integrity and Purpose Limitation, Access, and Recourse.

4.2 Seven Core Principles for AI in Education

The Center for Teaching Innovation outlines these guiding principles:

- The integrity of the faculty-student relationship
- A commitment to experimentation, evidence, and learning from experience
- Recognition that generative AI literacy is vital for students' post-Cornell careers
- Instructors have the duty to instruct and guide students on ethical and productive uses of AI
- Basic academic integrity principles remain important regardless of AI
- Integrating critique of AI practices and ethical issues into all stages of learning
- Detection algorithms are unreliable and not recommended for academic integrity enforcement

4.3 Flexible Faculty Framework

Cornell encourages a flexible framework where faculty and instructors can choose to:

- Prohibit generative AI use entirely
- Allow with attribution
- Encourage generative AI use

4.4 Research Findings on Learning Impact

Cornell's Center for Teaching Innovation cites emerging research showing mixed results. Some studies demonstrated learning gains in math and computer science, but often those gains proved ephemeral. Studies of AI's impact on writing are more concerning: students who used AI

for writing showed less brain activity, more homogeneity in their writing, and were less able to recall their written work.

5. University of Pennsylvania (Wharton)

Source: *Ethan Mollick's published frameworks and Wharton Interactive teaching resources*

5.1 Approach

Unlike most universities that set restrictive defaults, Wharton's approach under Professor Ethan Mollick (Co-Director of the Generative AI Labs) treats AI as a required learning tool, not an optional one.

5.2 Seven AI Roles in the Classroom

Mollick and Mollick (2023) define seven distinct roles AI can play in education, each with specific pedagogical benefits and risk profiles:

- **AI-Tutor:** Provides direct instruction and practice
- **AI-Coach:** Offers guidance and feedback on student performance
- **AI-Mentor:** Provides ongoing support and career guidance
- **AI-Teammate:** Collaborates on projects
- **AI-Tool:** Assists with specific tasks (research, writing, coding)
- **AI-Simulator:** Creates practice environments and scenarios
- **AI-Student:** Students teach the AI, reinforcing their own understanding

5.3 Key Policy Principles

- Students are required to use AI in assignments and then reflect critically on that use
- All AI use must be transparent and documented
- Students must remain the "human in the loop" with active oversight
- Critical assessment of AI outputs is mandatory -- students must evaluate for bias, errors, and accuracy
- Instructors are positioned as "builders and innovators" who design AI-powered learning experiences

5.4 Citation Requirements

Stanford GSB (which follows related principles) requires students to cite all resources including generative AI tools. Example citation format: "Gemini was used to gather sources for this project and create an initial written draft based on the author's outline. Substantial revisions were made by the author."

6. Massachusetts Institute of Technology (MIT)

Sources: raise.mit.edu; MIT AI & Education Summit 2025 proceedings

6.1 Institutional Position

MIT's Information Systems and Technology guidance advises students and staff to consider information security and data privacy when using AI tools. MIT does not have a single centralized AI policy; departments and instructors set their own guidelines.

6.2 RAISE Initiative

MIT's most distinctive contribution is the RAISE (Responsible AI for Social Empowerment and Education) initiative, which provides open-source AI literacy curricula used in over 170 countries. Key principles include:

- AI literacy is essential for all students, not just computer science majors
- Students should understand how AI works, its potential for good and harm
- Cultural representation matters -- AI trained on limited data can misrepresent or erase communities
- Open-source, multilingual learning tools should be prioritized
- Ethical standards should be universal in principle and flexible in practice

6.3 Key Research Findings

MIT's 2025 AI & Education Summit highlighted: the widening gap between AI tool availability and teacher readiness; the need to go beyond Bloom's taxonomy and ask students to critically examine and co-create with AI; and that access is not just about devices but also about language, trust, and cultural fit.

7. Duke University

Source: ctl.duke.edu/ai-and-teaching-at-duke-2/artificial-intelligence-policies

7.1 Faculty-Led Policy Framework

Duke's Center for Teaching and Learning encourages faculty to thoughtfully determine their stance, recognizing that standardized, one-size-fits-all policies are neither sustainable nor reflective of the diversity of teaching approaches.

7.2 Four Base Approaches

Duke presents four basic approaches instructors can take:

- **Full prohibition:** Students are not allowed to use advanced automated tools on assignments. Each student is expected to complete each assignment without substantive assistance from others, including automated tools.
- **Permission required:** Students are allowed to use advanced automated tools on assignments if instructor permission is obtained in advance.
- **Allowed with attribution:** Students may use AI tools but must acknowledge and cite all AI-generated content.
- **Fully encouraged:** Students are encouraged to explore and use AI tools as part of their learning process.

7.3 Transparency Requirements

Duke emphasizes explaining the rationale behind AI policies to students. Recommended reflection approaches include:

- Submitting a transcript of the AI conversation as an appendix
- Writing a reflective companion piece answering: What was your prompt? Did you revise the output? Did you ask follow-up questions? What did you learn?

7.4 Guiding Principles

Duke promotes two dimensions for evaluating AI use:

- **Cognitive dimension:** Working with AI should not reduce your ability to think clearly. AI should facilitate, not hinder, learning.
- **Ethical dimension:** Students using AI should be transparent about their use and ensure it aligns with academic integrity.

8. University of Sydney -- Two-Lane Assessment Model

Source: *Danny Liu (2023), University of Sydney*

8.1 The Two-Lane Model

Danny Liu at the University of Sydney proposes splitting assessments into two parallel tracks:

- **Lane 1 -- Verify knowledge (AI-free):** Proctored tests, oral exams, in-class work with no AI access. Designed to confirm that students have actually internalized the material.
- **Lane 2 -- Demonstrate AI competence:** Assignments where students demonstrate how they use AI critically and thoughtfully within their discipline. Evaluates AI literacy and judgment.

8.2 Rationale

The core argument is that trying to outrun or "outdesign" AI with increasingly complex assignments is unsustainable. Both lanes serve learning, and designing a course with a mix of each provides the most robust approach to maintaining academic integrity while building future-relevant skills.

9. Comparative Summary

University	Default AI Stance	Faculty Autonomy	Detection Tools	Data Policy	Unique Feature
ASU	Faculty decides per course	High	Not recommended	No sensitive data in public tools	Universal ChatGPT Edu access
Stanford	Treated as help from another person	High	Unreliable; not recommended	No confidential data	AI Playground sandbox
Harvard	Faculty decides (3 templates)	High	Discouraged	Level 2+ data prohibited	Respondus lockdown for exams
Cornell	Faculty decides (3 options)	High	Discouraged	No confidential data	7 core principles framework
Wharton	AI use required and reflected upon	Moderate	Not emphasized	Standard	7 AI roles pedagogy
MIT	Department-level	High	Not specified	Standard	RAISE open-source curricula
Duke	Faculty decides (4 options)	High	Not specified	Standard	Cognitive + ethical dimensions
U. Sydney	Two-lane model	High	Not specified	Standard	AI-free + AI-competence tracks

10. Recommendations for Adaptation

Based on this comparative analysis, the following principles emerge as consensus across leading institutions. These are particularly relevant for universities in developing countries or those establishing AI policies for the first time.

10.1 Universal Principles

- **Faculty autonomy within institutional guardrails:** Every university studied allows individual instructors to set course-specific AI policies within a university-wide ethical framework. No leading institution has imposed a blanket ban.
- **Transparency over detection:** Leading institutions universally discourage reliance on AI detection tools. Instead, they emphasize student disclosure, citation, and reflection on AI use.
- **Data protection is non-negotiable:** Every institution prohibits entering confidential, personal, or institutional data into public AI tools.
- **Process over product:** Assessment is shifting from grading final outputs to evaluating the learning process, including drafts, reflections, oral defenses, and AI interaction logs.
- **AI literacy as a learning outcome:** Understanding how AI works, its limitations, and its ethical implications is increasingly treated as a core competency.

10.2 Considerations for Haitian Universities

- **Cultural context:** AI systems are trained predominantly on English-language data. Policies should encourage students to critically examine AI outputs for cultural relevance, accuracy in Haitian and French-language contexts, and potential erasure of local knowledge.
- **Oral assessment tradition:** The French academic tradition's emphasis on oral examination aligns naturally with AI-resistant assessment strategies. This is a strength to build upon, not a limitation.
- **Access equity:** Consider providing institutional access to AI tools rather than assuming students can afford individual subscriptions. ASU's model of universal free access is instructive.
- **Start with principles, not rules:** Cornell's seven-principle framework and Duke's cognitive/ethical dimension model provide adaptable starting points that don't require constant updating as technology changes.
- **Build local AI literacy:** MIT's free, open-source RAISE curricula (available at raise.mit.edu) can be adapted for Haitian Creole and French contexts.

10.3 Source Links

- ASU Digital Trust Guidelines: ai.asu.edu/digital-trust-guidelines
- Stanford AI Policy Guidance: communitystandards.stanford.edu/generative-ai-policy-guidance

- Harvard Provost Guidelines: provost.harvard.edu/guidelines-using-chatgpt-and-other-generative-ai-tools-harvard
- Harvard FAS Guidance: oue.fas.harvard.edu/faculty-resources/generative-ai-guidance
- Cornell IT Guidelines: it.cornell.edu/ai/ai-guidelines
- Cornell CTI Principles: teaching.cornell.edu/generative-artificial-intelligence
- Wharton Teaching with AI: interactive.wharton.upenn.edu/teaching-with-ai
- MIT RAISE: raise.mit.edu
- Duke CTL AI Policies: ctl.duke.edu/ai-and-teaching-at-duke-2/artificial-intelligence-policies
- Thesify Top 20 University Policy Comparison: thesify.ai/blog/gen-ai-policies-update-2025