# Governing Generative AI at Stanford:

Recommendations for the use of

Artificial Intelligence by Undergraduates

and Faculty

**by**

**Sergio Escobar Repullo**

**Silvia Josefina Lombardo**

**Elena Kim**

**Jenin Al Shalabi**

**Stanford University**

**December 19, 2025**

We examine how artificial intelligence (AI) is currently used by stakeholders across Stanford University (Stanford) and evaluates whether existing institutional policies reflect the current state of AI usage on campus. Based on an analysis of Stanford's current AI policies, a review of peer institutions AI policies, and 20 interviews with Stanford stakeholders, we recommend Stanford consider the widespread use of AI when designing its institutional policies.

This process requires moving away from restrictive or ambiguous AI policies, which are likely ineffective and risk leaving students unprepared for an AI-driven world. Instead, we recommend promoting responsible, transparent, and informed use of AI so the Stanford community develops the necessary skills to engage with these technologies, by protecting sensitive data, mandating disclosure of AI use, conducting accuracy and bias checks, and preserving strong faculty autonomy over course-level rules.

## I. METHODS & LIMITATIONS

Drawing on a course-based qualitative project conducted for policy analysis and educational purposes, we held 20 semi-structured interviews with Stanford-affiliated stakeholders and external contacts using voluntary participation and de-identification (pseudonyms; no direct quotes attributed to identifiable individuals). Because this was a small, non-random sample designed to surface governance-relevant themes rather than measure prevalence, findings should be interpreted as illustrative, not statistically generalizable.

## II. CURRENT STANFORD UNIVERSITY AI POLICY

As mentioned in the Generative AI policy Guidance adopted in 2023,[1] Stanford should provide clear and updated policy guidance to students about AI usage that violates the Honor Code and offer instructors the necessary flexibility to explore how to appropriately use AI in coursework. The current default policy states that, absent an explicit statement from the instructor, using AI in

completing coursework is equivalent to receiving assistance from another person. As such, using AI tools to substantially complete any exams or assignments is a violation of the Honor Code.

---

[1] Stanford's Generative AI Policy Guidance can be found in the following link: https://communitystandards.stanford.edu/generative-ai-policy-guidance

For students, if there is any uncertainty about their instructor's AI use policy, they should seek prior clarification, as individual instructors may differ in their course-specific rules and departmental guidance. The current AI policy requires students to acknowledge and disclose the use of AI tools whenever there is doubt. For faculty members, there should be a clear obligation to specify whether and how students are allowed to use AI in their courses. Under the current policy, Stanford grants instructors both the freedom and responsibility to establish course-specific AI use policies in their syllabi.

Current guidance shows gaps between students and faculty. As both groups continue experimenting with emerging AI tools, Stanford should develop more practical, interim policy updates informed by surveys and community feedback.

## III. COMPARATIVE ANALYSIS ON THE USE OF AI FOR ACADEMIC PURPOSES

Although Stanford was among the first universities to develop an academic AI policy, other institutional experiences are also worth considering in the context of this Memo.

### A. COLUMBIA UNIVERSITY AI POLICY

Columbia University currently has a "work in progress" AI policy, reflecting the ongoing evolution of the law and the technology. The policy applies equally to students and faculty, a distinction that is central to this Memo. Its main features include: (i) contacting the IT center before procuring AI tools; (ii) prohibiting the input of confidential, personal, or IP-protected information; (iii) verifying the accuracy and potential bias of outputs; and (iv) disclosing any use of generative AI tools.

Under Columbia University's AI Policy, students may use AI tools only with explicit instructor authorization. Without such consent, using AI for assignments or exams is prohibited and considered a form of unauthorized assistance or plagiarism. With respect to instructors, Columbia's AI policy provides that faculties (i) should share clear expectations on the appropriate use of AI; (ii) should develop a policy course about the appropriate and inappropriate use of AI; and (iii) AI detection tools must be carefully used, preventing misidentification risks[2].

### B. HARVARD UNIVERSITY AI POLICY

Harvard University has issued "Guidelines for the Use of Generative AI Tools"[3]. According to these guidelines, Harvard contemplates the creation of local academic and administrative policies that complement the general framework described in the memo.

---

[2] Columbia's AI policy can be found in the following link: https://provost.columbia.edu/content/office-senior-vice-provost/ai-policy

[3] Harvard's AI guidelines can be found in the following link: https://www.huit.harvard.edu/ai/guidelines

The main provisions of Harvard's policy include: (i) protecting confidential data and prohibiting the entry of such information into publicly available generative AI tools; (ii) reviewing all AI-generated content before publication or distribution; (iii) complying with local academic and administrative policies; (iv) awareness against phishing and other cybersecurity risks; and (v) using only Harvard-approved AI tools for university-related work.

### C. Relevant literature

Regarding the use of AI for educational purposes, scholars have suggested key recommendations: (i) faculty members have the autonomy to determine the extent to which AI may be used, and they should establish clear policies and expectations; (ii) the use of AI should be disclosed, specifying how it was used; (iii) encourage teaching methods that frame AI as a starting point rather than a substitute; (iv) design assignments that minimize the temptation of AI-assisted ghostwriting; and (v) support AI-use policies with sanctions that are proportionate to the violation.[4]

## IV. Empirical evidence on the use of AI

To assess AI's impact at Stanford and identify perceived benefits and risks in academic settings, we conducted twenty 30-minute interviews in person, by videoconference, and by phone with key stakeholders, including students, faculty, parents, industry representatives, and administrators. All participants reported using AI, with 60% daily, 25% weekly, and 10% often. To ensure candid responses, anonymity was granted, with pseudonyms used in this memo. Interviewee demographics (department, age, nationality, role) are summarized in Appendix B Table 1, and a synthesized overview of each stakeholder's perceived benefits and risks appears in Appendix B Table 2.

### A. Key Takeaways

All stakeholder groups agree that AI is now deeply embedded in daily life and any policy must reflect that reality. They view AI's main benefits as greater efficiency, access to information, learning support, and research enhancement, while highlighting risks such as weakened critical thinking, inaccuracy, privacy and IP concerns, bias, and academic integrity issues. These topics provide the foundation for the stakeholder-specific findings that follow.

[4] Lemley, M. A., & Ouellette, L. L. (2025). *Plagiarism, copyright, and AI* [Manuscript in preparation]. *University of Chicago Law Review Online* (forthcoming).

## B. STUDENTS

a. Usage

All interviewed students, from undergraduates to post-doctoral fellows, reported using ChatGPT, with Will and Ethan paying for a subscription. Many also use Claude, Gemini, Copilot, Clauster, and other AI tools. Use cases vary widely. Charlie and Guido described using AI for "everything". Students cited uses from medical and lifestyle advice to academic clarification. In academic contexts, students commonly use AI to summarize readings and lectures, clarify complex ideas, brainstorm essay outlines, and correct grammar. Nine of the twelve students have used AI for coding, with Zoe, Claire, and Erik emphasizing its value for debugging and code completion.

b. Benefits

When asked about the benefits of AI, all interviewed students noted that it saves time by making tasks more efficient. Erik supported this by stating that AI "[expands] learning capacity." Charlie described AI as a private professor or study partner; Oliver referred to it as a technological team; and Ethan characterized it as a personal assistant. Several students, including Ariana, Will, and Alex, highlighted AI's role in writing and organizing ideas, while Alex and Claire noted its ability to support multilingual writing. AI's research value was repeatedly emphasized: Zoe called it a "great research tool," and Oliver stated that it "improves research productivity." Together, these descriptions illustrate how students increasingly view AI as an integrated learning resource.

c. Risks

Students identified several key risks associated with AI use. The most common concern was the loss of critical thinking from overreliance on AI: Joseph warned that students might "fall behind the curve," Guido feared rising illiteracy, Zoe cautioned that dependence could block genuine learning, and Alex described it as an "outsourc[ing] of the brain." Additional risks included privacy and IP leakage (Ethan, Erik, and Claire), hallucinations and accuracy issues (Claire, Ethan, Zoe, and Charlie), and the environmental impact of AI's energy and water consumption (Oliver). Together, these concerns underscore the tension between AI's benefits and its potential to undermine core educational capacities.

## C. FACULTY

a. Usage

All three professors interviewed reported using AI tools, and all have used ChatGPT. Dr. Lily also uses DALL·E and Runway ML, while Dr. Robin works with Claude and API-based tools. In

contrast, Dr. Sarah relies mainly on traditional search engines like Google and uses AI only occasionally. Lily and Robin employ AI for research, project development, art, and administrative tasks such as drafting syllabi, recommendations, emails, and data analysis, whereas Dr. Hensler refuses to use AI for writing. This highlights the wide variation in faculty adoption and usage styles compared to students' more consistent immersion in AI tools.

b. Benefits

Each faculty member highlighted a different benefit of AI. Dr. Lily sees AI as a tool that expands creative experimentation while democratizing design for students without technical training. They also explained AI is valuable because it streamlines repetitive preparative work. Dr. Robin echoed a sentiment expressed by the student body, it saves time. Dr. Hensler expressed the main benefit of AI is using it as a research tool for a first approach.

c. Risks

Faculty identified risks around originality, bias, and accuracy. Dr. Lily warned that AI can erode originality, as AI-generated imagery often imitates existing styles without true interpretation or meaning, raising ethical concerns about authorship and intellectual property. Dr. Robin cautioned against treating AI "as a person," noting that tools like ChatGPT can create feedback loops that amplify user biases while presenting information as objective. Dr. Hensler expressed concern over only using AI to "find the answers," mentioning hallucinations, retrieval errors, and the deceptive appearance of accuracy. Collectively, these views reflect faculty fears that AI may weaken critical judgment while concealing bias and inaccuracy behind an illusion of authority.

**D. PARENTS**

a. Usage

Both parents, Florence and Haley, showed a daily use of AI and revealed both strong educational potential and significant risks. Florence primarily uses ChatGPT as a research and teaching aid. Haley uses ChatGPT and Copilot to help compose messages, support academic writing, answer questions, and break down complex topics such as journal articles and math problems.

b. Benefits

Parents view AI as a valuable tool for individualized learning and for helping students plan, outline, and organize their work. Florence finds it especially useful for structuring content, improving communication, and saving time when preparing materials, while Haley highlights AI's ability to simplify complex ideas, support academic work, and accelerate medical and scientific research. Both believe AI can also help teachers refine their teaching methods and course design.

c. Risks

Florence worries that AI discourages genuine learning and discovery, potentially fostering dishonesty given its rapid, unreflective growth. Haley fears AI could exploit human creativity without compensation, devalue original work, and pose risks such as hacking and security threats. Both stress the importance of using AI ethically and consciously to ensure it supports rather than replaces authentic learning.

## E. INDUSTRY

a. Usage

Both industry stakeholders, Nick and Kate, use AI extensively. Nick uses ChatGPT, Perplexity, and Gemini weekly and works with "frontier" models through APIs to optimize auto-graders and support work tasks. Kate uses Gemini-generated search results and ChatGPT daily for personal usage and relies on foundation models professionally. Neither uses AI directly for academic work. However, Nick evaluates how recent graduates apply AI in their portfolios.

b. Benefits

Both interviewees emphasized AI's efficiency benefits. Nick noted that it streamlines hiring processes by reducing administrative work, while Kate highlighted its role in boosting productivity through data labeling, faster text processing, and affordable medium-quality creative work. Together, these perspectives reflect a view of AI as a tool that streamlines operations and expands individuals' capacity to work.

c. Risks

Industry stakeholders expressed broad concerns about AI's risks, including ethical, economic, and societal harms. Nick warned that overreliance on AI could homogenize thinking and compromise ethical hiring, while Kate cautioned about biased or harmful content and the exploitation of human creativity without compensation. Both feared that powerful AI systems might displace emerging professionals and erode the human talent pipeline. Overall, they viewed AI as offering major efficiency gains but also introducing significant challenges.

## F. ADMINISTRATORS

a. Usage

Javier reported limited but growing use of AI within his administrative unit. His department maintains a chatbot linked to the Stanford Law School Student Handbook. Beyond this, AI adoption is minimal, though he noted that faculty vary significantly in their personal and instructional use of AI.

b. Benefits

Javier emphasized that AI can help students generate ideas, clarify how to approach assignments, and improve the organization and clarity of their writing. He also noted that faculty can use AI to design activities, such as negotiation exercises, and to review course materials for consistency or clarity. Overall, he sees AI as a potentially valuable academic support tool when used appropriately and together with clear expectations.

c. Risks

Javier's primary concern centers on academic integrity and legal risk. He highlighted copyright issues and emphasized that AI use in exams can conflict with the Honor Code, which prohibits AI unless explicitly allowed. He noted that take-home exams are now harder to administer because of the difficulty of controlling AI use.

## V. RECOMMENDATIONS FOR THE USE OF AI BY UNDERGRADUATES AND FACULTY

As a preliminary note, all recommendations in this Memo are based on the current state of AI. Should transformative developments occur, such as the emergence of Artificial General Intelligence, we strongly advise that Stanford's policies be revisited. That said, in the present context, we offer a set of recommendations for a new Stanford AI Policy, which we believe may be useful for the next five years.

As a general framework, we recommend the creation of a "Code of Practice for the Use of AI", establishing key principles that all members of the Stanford community must follow when engaging with AI tools. To ensure that these guidelines are inclusive and contextually appropriate, they should be developed collaboratively with relevant stakeholder groups.

Additionally, in the next sections this Memo sets forth some policy recommendations for students and faculty, based on the evidence collected and the differing academic roles and responsibilities of each group.

1) Transparency & Disclosure: Clear expectations for when/how students and instructors disclose AI assistance (with brief context, not just "yes/no").

2) Data Protection & Procurement: Rules for what data can/can't be entered into tools, plus vendor/integration controls (contracts, retention, logging, third-party access).

3) Academic Integrity & Assessment Design: Course-level guidance that reduces incentives for ghostwriting (process-based assignments, oral checks, drafts, scaffolded work) and defines unacceptable uses.

4) Oversight, Auditing & Iterative Updates: A lightweight mechanism to track incidents, update

guidance on a fixed cycle (e.g., annually/each quarter), and publish summary metrics and revisions.

### B. RECOMMENDATIONS FOR THE USE OF AI BY UNDERGRADUATES

From the evidence collected for this Memo, the most common risks identified for students' use of AI were (i) the loss of critical thinking, (ii) risks to academic integrity, and (iii) AI-related errors such as hallucinations. Therefore, the following measures are suggested to address them:

1) Students are required to disclose when AI was used and explain how it was used. Each department must determine how they intend to promote the use of AI. For example, law schools may promote AI use for augmentation purposes (e.g., when drafting papers), while physics departments may choose to allow automation purposes (e.g., using AI for coding).

2) Promote assignments and activities where students are not allowed to use AI or computers. The purpose of this measure is to foster students' critical thinking skills by engaging them in tasks that do not rely on technology.

3) For writing assignments, instructors are encouraged to hold follow-up meetings with students to verbally discuss their work.

4) Need to reevaluate the benefits of take-home exams, opting instead on in-class exams or the use of Exam4 software to ensure independent performance.

### C. RECOMMENDATIONS FOR THE USE OF AI BY FACULTY

According to the evidence gathered for this Memo, the main concerns regarding the use of AI by faculties are that (i) they should have the freedom to determine how to use AI, and (ii) they must develop the skills necessary to teach students the responsible and effective use of AI.

1) Stanford's AI Policy should explicitly preserve faculty discretion to determine whether and how students may use AI tools in their coursework.

2) Faculty members must also be knowledgeable about how to use AI responsibly, in order to guide students in its proper and ethical use.

3) Stanford's AI Policy should encourage the recruitment of experts in AI and education and promote faculty training on the responsible use of AI. This could include periodic seminars or courses led by professors specializing in AI and education, designed to help faculty stay informed about emerging trends and best practices in the field.

### D. RECOMMENDATIONS FOR THE ENFORCEMENT OF THE POLICY

Considering our proposal that faculties retain autonomy to determine the use of AI use in their courses, we suggest the following measures to align enforcement with that policy:

1) If a faculty member prohibits AI use for a specific task, the policy should clearly define the corresponding sanctions and procedures, ensuring that both students and faculty understand the consequences in advance.

2) To ensure consistency and fairness, we recommend that Stanford AI policy establish a committee responsible for reviewing violations of the AI-use policy and recommending appropriate disciplinary actions.
3) The AI policy should include a proposal to amend the Honor Code by adding an explicit "AI clause." This would ensure that the unauthorized use of AI is treated with the same seriousness as other forms of academic misconduct.

4) Finally, any sanction should be proportionate to the severity of the violation. For example, a student who uses AI to generate and submit an entire paper without review (automation use) should face a different level of penalty than one who merely uses AI for grammar correction.

### E. Practical elements for the Enforcement

To translate these recommendations into practice, Stanford's Office of the Vice Provost for Undergraduate Education could pilot a two-tier AI Governance Framework: (i) a short-term AI Policy Hub, coordinating with coordinate with the Center for Teaching and Learning and the Office of Community Standards to centralize guidelines and disclosure practices while preserving departmental flexibility, and (ii) a long-term AI and Academic Integrity Task Force to monitor effectiveness, assess new technologies, and recommend annual policy updates.

We recommend Stanford launch a "Transparent AI" campaign that frames disclosure as a marker of academic integrity and digital literacy, and not as punishment. Integrating AI governance into Stanford's educational mission would promote an ethical use of AI. A preliminary overview of stakeholder roles and considerations is included in Appendix A.

## VI. Conclusion

Stanford's challenge is not to eliminate AI from the classroom, but to teach students and faculty to use it intelligently, transparently, and responsibly. The policies outlined in this memo embrace that balance: autonomy with accountability, innovation with integrity. As AI continues to shape the conditions of knowledge itself, Stanford can define what ethical engagement looks like in practice by turning uncertainty into a new standard for excellence.

**Word Count**: 2994 words.

**DISCLOSURE OF AI:** *AI Technology Used:* ChatGPT. *AI Technology Usage:* Grammar check; make paragraphs more concise and clearer; clarification of ideas; and synthesized survey data for Appendix Table 2. All the uses were performed with human supervision.

## APPENDIX A

| Overview of Stakeholder Roles and Considerations | | | |
|---|---|---|---|
| **Stakeholder** | **Role in Implementation** | **Primary Benefits** | **Potential Concerns** |
| Students | Follow disclosure requirements; engage in transparent and responsible use of AI tools | Clear expectations, academic integrity support, skill development in ethical AI use | Fear of punitive enforcement; uncertainty about allowed tools |
| Faculty | Create and publish course-specific AI policies; integrate ethical AI pedagogy | Consistency across courses; guidance and flexibility maintained | Additional workload; concerns about autonomy in classroom governance |
| Office of Community Standards (OCS) | Provide infrastructure for disclosure reporting and adjudicate violations | Stronger foundations for fair enforcement and precedent-setting | Increased processing burden with new cases |
| VPUE (Vice Provost for Undergraduate Education) | Central coordination of AI governance; oversees annual policy review and communication | Demonstrated leadership on national AI ethics; improved campus alignment | Requires continuous resourcing and cross-unit collaboration |
| CTL (Center for Teaching & Learning) | Develop faculty training, templates, and instructional materials | Scalable support; more efficient rollout | Needs staff capacity and time to build resources |
| Technology Partners (e.g., University IT) | Support tooling for disclosure tracking and responsible AI access | Clarity and usability for students/faculty. data-informed policy refinement | Implementation costs and maintenance |

## APPENDIX B

Table 1. Interviewee Demographics

| ID | Name | Role | Department / Organization | Age | Nationality |
|---|---|---|---|---|---|
| 1 | Ariana | Undergraduate Student | Political Science | 20 | United States |
| 2 | Will | LLM Student | Law School | 28 | - |
| 3 | Charlie | LLM Student | Law School | 35 | - |
| 4 | Joseph | Master Student | Engineering | 22 | French |
| 5 | Guido | Master Student | Business | 29 | Argentina |
| 6 | Oliver | Master Student | Civil & Environmental Engineering | 24 | United States |
| 7 | Erik | Master Student | Computer Science | 27 | Indian |
| 8 | Zoe | Master Student | Aeronautics & Astronautics | 24 | - |
| 9 | Ethan | JD Student | Law School | - | - |
| 10 | Leo | PhD Student | Political Communication | 30-35 | Brazil |
| 11 | Alex | Post-Doctoral Fellow | Sociology | 30-35 | South Korea |
| 12 | Claire | Post-Doctoral Fellow | Medicine & Health Policy | 40 | Lebanese |
| 13 | Lily | Associate Professor | Theater & Performance Studies | 32 | American |
| 14 | Robin | Professor | Sociology | 35-40 | Taiwan |
| 15 | Sarah | Professor | Law School | 82 | United States |
| 16 | Florence | Parent | - | 61 | Spain |
| 17 | Haley | Parent | - | 61 | United States |
| 18 | Nick | Industry | Senior Talent Manager | 38 | Singaporean |
| 19 | Kate | Industry | Apple Manager | 35-40 | United States |
| 20 | Javier | Administrator | Law School Registrar | 74 | Peru |

| Table 2. Summary of Benefits & Risks per Interviewee | | | |
|---|---|---|---|
| **ID** | **Name** | **Benefits** | **Risks** |
| 1 | Ariana | Saves time on readings; improves writing flow; helps organize ideas. | Over-reliance weakens independent writing; unclear disclosure rules create anxiety. |
| 2 | Will | Increases efficiency; helps structure and deepen ideas; supports creation of polished deliverables. | Reduces critical thinking; fosters dependency on AI systems. |
| 3 | Charlie | Acts as a "private professor" or study partner; improves efficiency, especially for readings. | Hallucinations; risk of academic dishonesty when AI-generated ideas are not disclosed. |
| 4 | Joseph | Faster, convenient tool for testing and exploration. | "Falling behind the curve"; decreased self-intelligence; environmental concerns. |
| 5 | Guido | Saves time while maintaining human-level quality of work. | Overuse could increase illiteracy, and weaken reading and writing skills. |
| 6 | Oliver | Gives laypeople access to technical skills; improves research productivity; functions as a "technological team." | Climate impacts (water use, energy demand); negates efficiency gains; technological unemployment; consumer obsolescence; "alarming risks." |
| 7 | Erik | Automates repetitive coding; speeds experimentation; expands learning capacity. | Overfitting to AI-suggested patterns; privacy and IP leakage. |
| 8 | Zoe | Helps them be a better student; supports studying, intuition-building, efficiency; strong research aid. | Dependency; lack of genuine learning; hallucinations; false information; AI's "always affirming" behavior; risks of social isolation and parasocial attachment. |
| 9 | Ethan | Functions as a personal assistant; provides feedback; increases time efficiency. | Hallucinations; data privacy risks; harms judgment and critical thinking if used from the start of a task. |
| 10 | Leo | Saves time. | Information overload; time wasted; fear of over-reliance on AI. |
| 11 | Alex | Saves time; supports writing in English; helps with phrasing and tone. | Deprives users of learning; "outsourcing the brain"; less personal ownership over work; increased self-doubt. |

| | | | |
|---|---|---|---|
| 12 | Claire | Saves time; improves accessibility; supports multilingual researchers. | Potential bias in medical phrasing; privacy concerns; accuracy issues. |
| 13 | Lily | Expands creative experimentation; democratizes design; streamlines repetitive prep work. | Erodes originality; AI imitates but cannot interpret; ethical ambiguity around authorship and IP. |
| 14 | Robin | Saves time; has become hard to imagine working without AI. | Tendency to treat AI as a person; bias amplification; illusion of objectivity; "living inside your own ideas." |
| 15 | Sarah | Strong research tool for finding sources and exploring answers. | Hallucinations; retrieval errors; subtly wrong but convincing information; difficulty seeing what's missing. |
| 16 | Florence | Organizes content; fine-tunes structure; supports communication and planning; saves time. | Undermines thinking and discovery; may encourage dishonest behavior; society has not fully reflected on AI's rapid growth. |
| 17 | Haley | Helpful for learning; advances medical research and science; supports rapid development. | Steals creative output without compensation; devalues human creativity; unknown risks; hacking and security threats. |
| 18 | Nick | Boosts efficiency; reduces biased language; saves recruiting time. | Homogenizes thinking; over-reliance reduces innovation; ethical issues in automated hiring. |
| 19 | Kate | Enhances productivity; supplements human labeling; helps with spelling; enables cheaper medium-quality creative output; can support learning. | Direct harms (harmful/offensive content, bias, stereotypes); exploitation of human creators' data; crowding out in-training workers, weakening human expertise and future training data. |
| 20 | Javier | Generates ideas for students; helps them structure and improve writing; supports faculty in designing activities and checking materials. | Copyright concerns; honor code violations in exams; pushes instructors toward closed-book exams due to fear of AI misuse. |