OFFICE OF THE PROVOST
UNIVERSITY OF MICHIGAN

INFORMATION AND TECHNOLOGY SERVICES
UNIVERSITY OF MICHIGAN

AI in Education Working Group

April 30, 2026

I. Executive Summary

This report evaluates the status of artificial intelligence (AI) in education at the University of Michigan at the moment and offers recommendations for the strategic direction for the near and medium-term future. It builds on prior work at the university and acknowledges the fast pace of changes, which necessarily make all recommendations speculative.

In simplest terms, the challenge of AI for education is the following: **AI systems have been developed to speed up processes; learning requires processes that slow things down.** To learn is to confront *desirable difficulties* and overcome them through *productive struggle*.[1] AI systems are already increasing efficiency and economies of scale; they make possible efforts that may have been literally impossible only a few years earlier. Much of this is useful at universities, especially in their research missions and business operations. Even there, AI tools, like all new technologies, have their potential costs. But it is in education that AI's fundamental purpose to make things easier and faster poses a particular challenge.

The challenge of AI brings to higher relief the essential purpose of university education in the twenty-first century: **to foster competencies for human agency and judgment**. Even in a technologically mediated future, university-educated people will have roles in which humans bear the final responsibility for actions. Their ability to make sound judgments, including about the uses of technology, is what universities prepare them for.

Informed by that general principle, this report recommends the following:

- Develop **instructional design resources and AI literacy programming** to help instructors decide whether, when, and how to integrate AI tools confidently and critically.
- Develop **AI literacy outcomes for all its graduates**. Units should be invited to craft these in discipline-appropriate ways. This report outlines, at a general level, what those competencies should include.
- Recognize that **AI integration does not mean universal AI adoption**.
- **Avoid surveillance-based AI detection tools** to address assessment challenges, but instead revisit potentially outdated practices around assessment and **invest in a broader assessment infrastructure**.
- Set up **ongoing advisory bodies** consisting of faculty, staff, and students to inform university decisions about AI.

[1] Bjork, Robert A., and Elizabeth L. Bjork. "Desirable Difficulties in Theory and Practice." *Journal of Applied Research in Memory and Cognition* 9, no. 4 (2020): 475–79. 2020-99542-008. https://doi.org/10.1016/j.jarmac.2020.09.003.

- Invest in the **technical infrastructure** necessary for experimentation with pedagogical and technical approaches.

Recognizing that U-M's decentralization is both a virtue and a challenge, the working group recommends a process where the university, as an institution, articulates a strategic vision based on the ideas above and empowers units to develop their own guidelines, practices, and, where necessary, policies around how that vision is best realized, given unit missions and disciplinary conventions.

Finally, the working group believes **action on AI is urgent**. First, the University of Michigan has an opportunity to help set the direction for how AI is integrated into higher education, on our terms instead of the terms of external forces. Second, as teachers, many U-M faculty feel a loss of control in their own teaching. This report urges the university to address both elements as promptly as feasible.

Table of Contents

II. Background and Charge

The AI in Education Working Group was established by its co-executive sponsors, Provost Laurie McCauley and VPIT-CIO Ravi Pendse. It was charged to explore the thoughtful, ethical, and effective integration of artificial intelligence into teaching and learning at the University of Michigan. This exploration included the relationship of AI to models of liberal arts inquiry and instruction, as well as circumstances in which AI may not be appropriate for use based on desired learning outcomes. The working group was charged to consider balancing AI and Large Language Model (LLM) usage with maintaining core knowledge in the fundamentals of the discipline of study and promoting critical thinking skills.

The group built on prior working group reports by reviewing their findings, assessing current initiatives, and identifying future opportunities to strengthen the educational impact of AI across the university. These reports included, most importantly, the June 2023 Generative Artificial Intelligence Advisory Committee Report (known as the GAIA Committee Report) as well as reports on high-performance computing and AI in research at Michigan.

The working group's focus on AI in education meant two related but distinct sets of questions:

1. **The opportunities and challenges presented by AI *for* education**. As elaborated in the sections below, this included both pedagogical approaches as well as assessment and academic integrity across the university units and student bodies.
2. **Education *for* AI**. The working group was charged to articulate (at a general level) the AI competency U-M graduates should possess.

Given this focus, this report does not address AI or LLMs in either research or for university business uses. There are some obvious overlapping areas: most faculty and students have at least some involvement in research activities, and university administrative tools include learning management systems. Inasmuch as these are considered in this report, it is always from the education perspective.

III. Approach and Process

The working group consisted of thirteen members with faculty appointments and two staff, one from the Provost's Office and one from ITS. Members included faculty members from the Ann Arbor, Dearborn, and Flint campuses as well as Michigan Medicine. Their expertise ranged from the arts and humanities to social sciences and STEM fields. The three largest colleges on the Ann Arbor campus (LSA, College of Engineering, the Ross School of Business) were represented as well as a number of professional schools. Members' perspectives, expertise, and ways of using AI systems ranged from non-users to active developers of systems. The membership is listed in Appendix A.

The working group began its meetings on December 19, 2025, and met roughly once every two weeks through winter 2026, primarily in a hybrid format. The group began by reviewing prior U-M reports, identified outstanding questions and landscape changes, and set the distinction between AI *in* education (challenges and opportunities) and education *for* AI (whether and what competencies the university should provide to its students) as its main categories of focus. The working group scheduled listening sessions for individuals who had raised concerns about or were known for innovative uses of AI in teaching, and people who had found solutions to the challenges AI poses. It also commissioned a survey of all university faculty on AI, received results of a survey done of undergraduates in fall 2025, and commissioned a benchmarking study of AI in the U.S. higher education. In late winter term, the working group co-sponsored a survey of all university graduate students, but results were not available by report deadline.

The working group or its individual members met with different stakeholders for informational conversations or consultations, collected information from outside the university, and received and reviewed unsolicited information.

Working group membership — except the chair's identity — was kept confidential at member request during its work. Internal conversations were confidential to allow for a frank exchange of views. Recommendations in this report reflect a broad working group consensus. Where differences of opinion existed among members, they are discussed, but without identifying individuals.

In addition to the formal charge, the working group adopted a *practical orientation* early as an additional guiding principle. It was clear from initial conversations and confirmed by the survey and the listening sessions that people wish for practical guidance and concrete solutions. At the same time, the working group recognizes that solutions focused on specific tools or systems may become outdated very fast. This report aims at a balance between those two considerations.

The final report was written collectively by the entire working group. Different members drafted different sections, some in small groups and some alone. These drafts were then reviewed and discussed by the entire working group. The working group chose to retain the slight variation in style and tone between sections. Because of the overlapping themes, some observations and recommendations are reiterated in different parts.

The working group used generative AI tools in its processes. Different members used it for brainstorming, structuring information and ideas; some did not use it at all. The entire report has been edited by humans.

Representing Faculty Viewpoints and Strengths

The working group offers four observations about faculty positions on AI in the teaching and learning environment. A range of these positions was represented in the working group membership; as the faculty survey shows, the range exists at U-M generally.

- **AI skepticism is a valid intellectual position, not a deficit to overcome.** A framework that treats faculty resistance to AI as a low-information or a change-management problem will produce poorer results. A framework that treats faculty skepticism as a source of thought-provoking critique — to be engaged on its merits, and to sharpen the case for where AI integration is and *is not* appropriate — will produce stronger pedagogy across the board, including in the courses where AI is ultimately used heavily.
- **AI enthusiasm exists on continuums.** AI enthusiasm and skepticism exist on continuums; faculty who design AI systems think about dangers and limitations. Faculty who are skeptical of LLMs may be supportive of some other AI systems. Faculty and students who are less enthusiastic about AI may also be well-informed about how the technology works.
- **Supporting pedagogical experimentation — in many directions.** The university's reputation and its learning environment are fundamentally strengthened when students take highly experimental courses using new AI systems and technologies *and* in courses that make use of low or no-tech spaces. An approach that takes both seriously provides proof that our students are AI literate, develop cognitive skills such as focus, *and* are critical thinkers without technology aids. An approach that leans into AI education alongside low-technology spaces will positively separate U-M from peer institutions. Faculty experimenting in many directions need support from administrators, requiring protection from grievances rooted in comparison with traditional expectations, as well as infrastructure support.
- **Support for faculty innovators and peer learning rather than top-down training.** The university should support faculty innovators. Faculty innovators should share their pedagogical practices with colleagues in their own disciplines or schools to encourage new and effective learning communities in the AI environment. The institution should identify and resource these faculty — including with course release, travel funds, and recognition — rather than substituting centralized AI training for them.

IV. The Landscape of AI in Education

To understand the current state of AI in education, the working group commissioned a survey of U-M faculty, reviewed the results of a student survey conducted in fall 2025, and conducted a benchmarking study of AI in higher education. The working group also co-sponsored a survey of U-M graduate students, but results were not available by report deadline. Given the working group's focus on AI in education, staff were not surveyed for this report. The working group also invited a number of U-M faculty members for listening sessions. The individuals were selected on the basis of recommendations identifying them as colleagues with thoughtful ideas about AI in education. They included innovative uses of AI, creative solutions to pedagogical or assessment challenges, and critics of the university's approach to AI.

Faculty Survey

The faculty survey was commissioned from Professor Josh Pasek (LSA Communications and Media, LSA Political Science, Institute for Social Research). The survey was sent to all U-M faculty in Ann Arbor, including Michigan Medicine, as well as in Dearborn and Flint: tenure-stream faculty (assistant, associate, and full professors, emeriti and emeritae), LEO faculty, clinical, research, and adjunct faculty, as well as librarians. Overall response rate was 8%, or 963 responses. Responses were roughly proportional to both school or college-level units and faculty ranks, with a few significant exceptions: the College of LSA was significantly overrepresented in the responses, while Michigan Medicine and College of Pharmacy were underrepresented. Similarly, full professors and LEO faculty were overrepresented, while clinical faculty were underrepresented. Details on these issues, general methodology, and additional information are provided in Appendix F.

Summary

Initial analyses of the faculty survey indicate that U-M faculty views on AI range widely from considerable skepticism to enthusiasm, but are less polarized than public discourse might suggest. Faculty across the spectrum share some common concerns. Most faculty across the institution believe the following:

- The university should prepare its students for **critical**, **ethical** and **effective** engagement with AI.
- The university should not ban all AI use or encourage total student AI avoidance.
- AI presents significant challenges to teaching, particularly to the assessment of learning, but faculty who have adjusted their teaching to AI have not found it to reduce learning.
- Faculty perceive the **potential** harms of AI as greater than the potential benefits for **critical thinking** and **writing abilities**.

Faculty most concerned about AI feel the university's messaging about AI has been one-sided, excessively enthusiastic, and with inadequate attention to its negative aspects. Although these views appear to represent a small minority, the working group took them very seriously in its deliberations. A summary of these views is presented in Appendix F.

The overall results do not significantly vary by faculty field, unit, position, or time at the institution.

General Findings

Based on the imputed results from the faculty survey, here are some of the general findings regarding faculty awareness, usage, and pedagogical responses to artificial intelligence:[2]

Awareness and Adoption

1. **High Generative AI Adoption:** A vast majority of respondents (85%) reported that they have used generative AI, compared to only 15% who have not.
2. **Frequency of Generative AI Use:** Over half of the users integrate generative AI into their routine regularly; 15% use it "every day," 16% use it "most days," and 22% use it "one to three times per week".[3]
3. **Recent Onboarding to AI:** The majority of faculty started using generative AI relatively recently, with 41% of users (35% of all faculty) reporting they started "last year" and 41% of users starting "two to three years ago" (34% of all faculty).
4. **Self-Rated Effectiveness:** Many respondents feel competent at getting their desired outputs from generative AI. When asked, 42% of those who have ever used AI feel "somewhat effective," 26% feel "very effective," and 6% feel "extremely effective."
5. **General AI Knowledge:** When assessing their overall knowledge of AI, 48% claim to know a "moderate amount," 21% know "a lot," and 10% know "a great deal". Only 1.3% say they know "none at all."

Pedagogy and Course Administration

The questions generating the following responses were only administered to faculty who indicated they have a teaching role at the university.

6. **Vulnerability of Assignments to AI:** Teaching faculty recognize that generative AI could competently handle much of their coursework. When asked how successful AI would be at completing the majority of their assignments, 41% said "somewhat successful," 16% said "very successful," and 7% said "extremely successful". Only 12% think it would be "not at all successful."
7. **Time Spent "AI-Proofing":** The amount of time faculty spend modifying curricula to prevent students from using AI (AI-proofing) is highly varied: 23% spent "no time at all," 26% spent "a little time," and 26% spent a "moderate amount of time."
8. **Effectiveness of "AI-Proofing":** A large portion of faculty lack confidence in their ability to AI-proof their classes. When asked how effective they were at designing curricula that could not be completed by AI, 21% reported that they were “not at all effective,” with a further 33% reporting that they were "somewhat effective."
9. **Impact of AI-Adjustments on Learning:** When faculty adjusted their curriculum for AI, the largest group (38%) observed "no noticeable change" in the learning experience. Among those who did notice

[2] Given the number of questions, the survey was designed for and distributed as *planned missingness*: for most of the questions, respondents were randomly assigned a subset of questions to make the survey shorter. For these summary results, the percentages were imputed, i.e., calculated estimates on what the percentage would have been if all the respondents had had and answered the question.
[3] If we account for those who have never used it, 13% report using generative AI "every day," 14% use it "most days," and 29% use it "one to three times per week", constituting a total of 46% who are using it weekly or more often. In contrast, 4% of those who ever used AI reported that they now “never” use it, comprising 19% of all faculty when combined with those who have not ever used it.

changes, opinions were relatively split, with 21% reporting the experience was "slightly" or "significantly" improved and 16% reporting that it was "diminished."

10. **Instructor AI Use for Course Materials:** Faculty use of generative AI to specifically create or update course materials remains low. 48% of those who ever used generative AI reported they "never" use it for this purpose (67% including AI nonusers), and 19% use it "rarely" (12% including nonusers) while on 10% reported that they usually or always use it (8% including nonusers).

Perceived Harmfulness and Usefulness of AI for Students

The survey asked faculty to consider areas in which they think AI can be either harmful or useful for students.[4] The categories were not mutually exclusive — a person can perceive both harms and benefits — and respondents could select any that apply. The table below summarizes the responses.

Category	Harmful %	Helpful %
Critical thinking	95	17
Content knowledge	75	54
Research methods	56	47
Writing ability	88	44
Technical skills	47	34
Clinical skills	27	13
Professional skills	51	27
Artistic or creative skills	48	18

Although respondents perceive both potential harms and benefits, they regard potential harms as much greater than potential usefulness for **critical thinking**, **writing ability,** and **creative skills.**

University Responsibility

Faculty views at the university largely converge on many *prospective* questions regarding AI. Figure IV.1 displays faculty responses to questions about what the university should be responsible for. As the figure shows, regardless of faculty's self-identified field of expertise, they believe it is extremely important the university should teach students to use AI **ethically**; very important to teach students to use AI **effectively**, teach students about the **negative impacts of AI**, and **protect the privacy of student data from AI**. There is slight variation on whether the university should prepare its students to use AI for their future careers (with humanists' and artists' views varying from that of others'). Humanists and health scientists tended to agree more than other groups that the university should show students why to never use AI. Overall, however, total AI discouragement at the university level is considered as "unimportant" across all groups.

[4] "In which of these areas do you think student use of generative AI can be relatively harmful?" and "In which of these areas do you think student use of generative AI can be relatively helpful?"

Appendix F shows that these convergences remain also when respondents are disaggregated by their actual position, time as a faculty member at the U-M, or campus and Ann Arbor school or college.

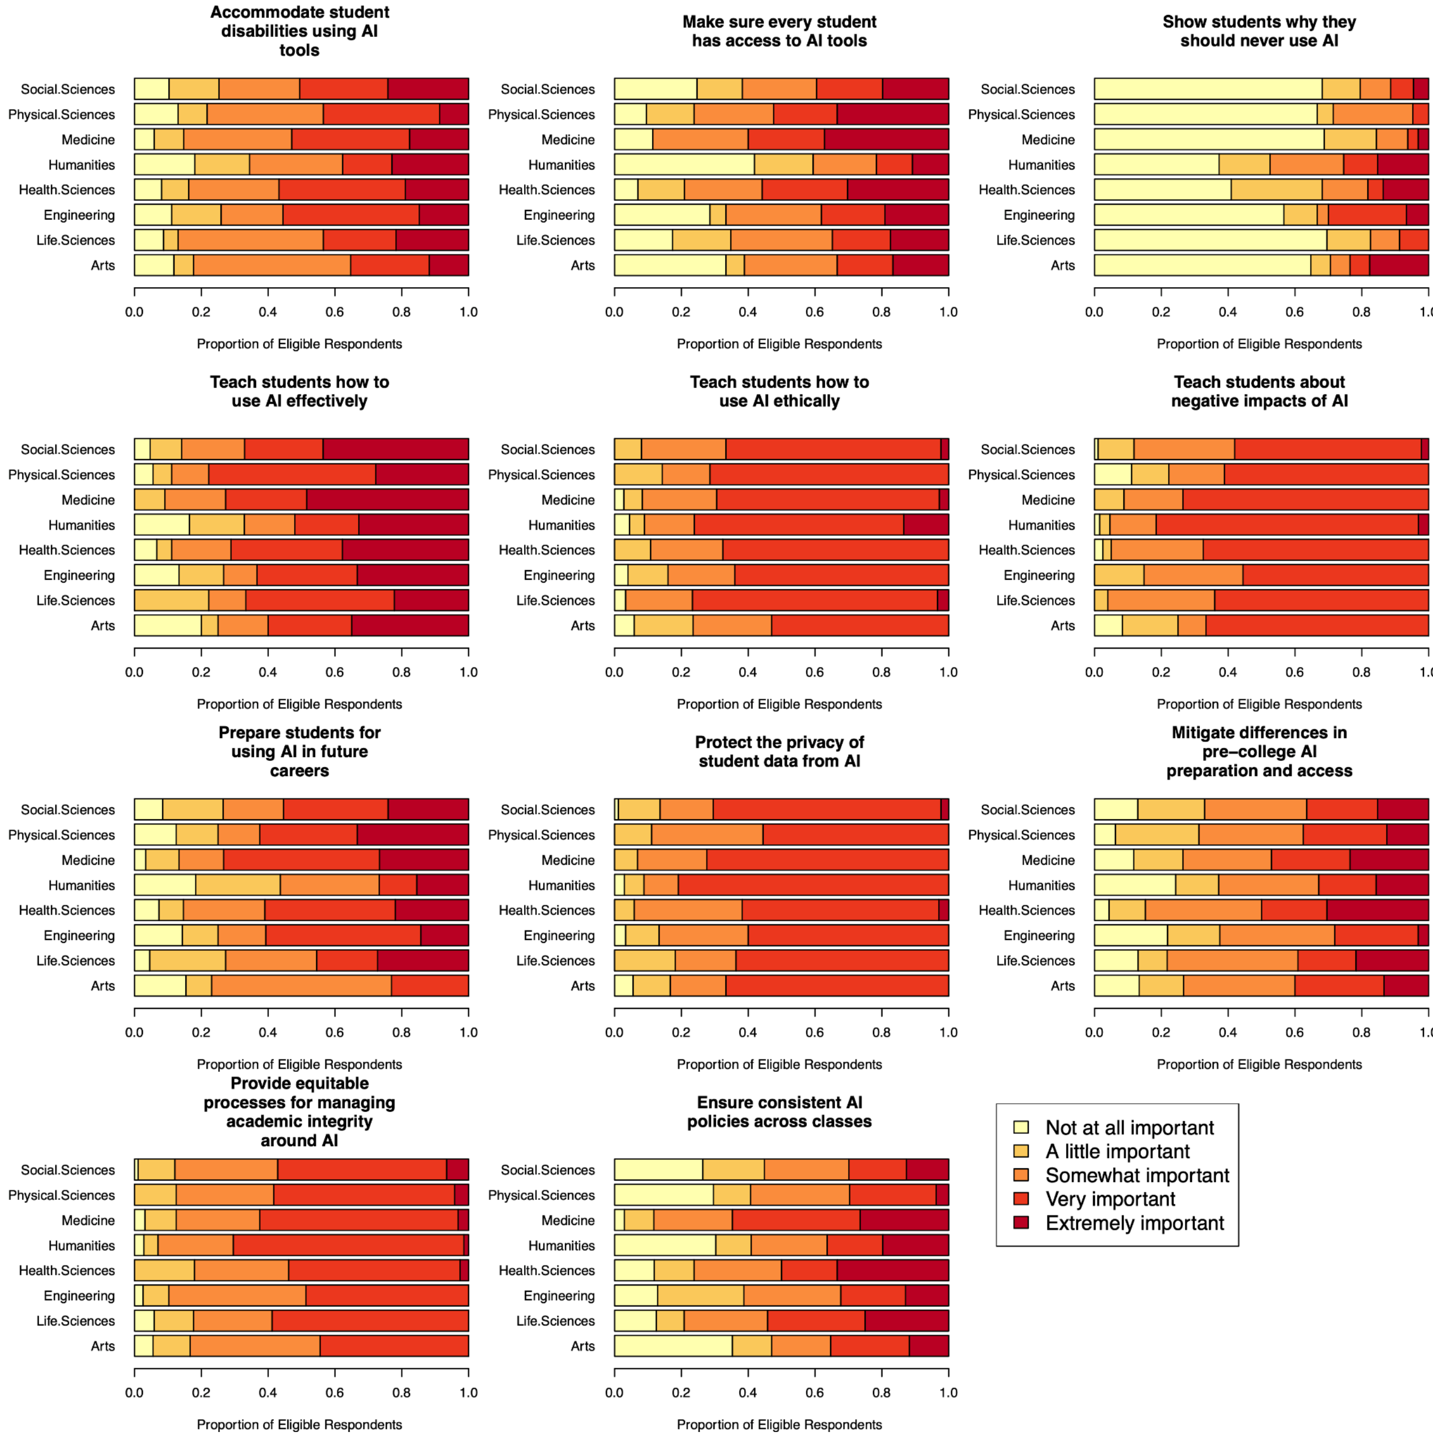

Figure IV.1: Faculty responses, disaggregated by self-reported field of expertise, on questions about what the university should be responsible for regarding AI.

Student Survey

An undergraduate AI use survey was conducted in fall 2025 in preparation for a Provost's Seminar on Teaching (PSOT). It had been commissioned by the Center for Research on Learning and Teaching, which organizes the PSOTs, and was conducted by Professor Pasek. Key findings were made available to the working group.

The survey was distributed by several instructors who volunteered to share it with the students in their courses. It therefore was not a formally representative sample. However, with 1,611 responses from all of the nineteen Ann Arbor schools and colleges it nevertheless offers a valuable snapshot, especially of undergraduates. Most respondents were first (**37%**) or fourth-year (**34%**) undergraduates. Only **3%** were graduate students. The majority of respondents were from **LSA (41%)**, **Ross (40%)**, and **Engineering (13%)**.

In terms of student self-reported knowledge, more than **50% of students said they had moderate AI knowledge**, and **31% cited having a lot or a great deal of AI knowledge**. Only **15% of participants had little to no AI understanding**. More than **45% reported having a moderate amount of experience with AI** and **31% a lot or a great deal**. The remaining **23%** of participants had little to no AI experience.

Student Key Uses: Information Seeking and Writing

Among all surveyed students **(n=1,611),** those who reported using AI for **writing**, **coding**, or **summarizing** were asked additional questions to better understand their motivations and use patterns.

- *Task Types and % of Respondents who use AI*
    - *Summarizing information (e.g., readings, notes, text, lectures, etc.)* ***(68%)***
    - *Searching for information/research* ***(68%)***
    - *Writing or editing text* ***(63%)***
    - *Solving quantitative or mathematical problems* ***(60%)***
    - *Planning for or outlining work tasks* ***(44%)***
    - *Personal or work life management (e.g., scheduling, routine tasks)* ***(41%)***
    - *Coding and programming assistance* ***(33%)***
    - *Generating creative content (e.g., images, video, audio)* ***(30%)***
    - *Something else* ***(13%)***

No Single Approach to Writing with AI

- **799** students (**63%**) use AI as part of their writing and editing process. These undergraduate students report using AI for a variety of tasks, including:
    - *Editing at the End*
        - **80%** of students use AI to proofread and copy edit their work at the end of the writing process.
    - *Brainstorming*
        - **77%** of students use AI to generate ideas and brainstorm.
    - *Turning Ideas into an Outline*
        - **66%** of students use AI to transform their ideas into an organized outline or structure.
    - *Producing Text*
        - **28%** of students use AI to generate a first draft and 14% have AI produce an entire text.

AI Writing Benefits: Feedback and Time Saving

- *Feedback and Efficiency*
    - Feedback (**72%**) and time-saving (**66%**) were the leading benefits reported by students using AI for writing and editing.

- *Idea Generation and Organization*
  - More than half of students reported using AI to support idea generation, organization, and structural planning.
- *Confidence and Accessibility*
  - A smaller share indicated that AI helps address writing anxiety (**19%**), language barriers (**17%**), and creative challenges such as generating new ideas (**40%**).

AI as a Study Tool

- Among the **871** students (**68%**) who reported using AI to summarize information were asked how they apply it across different academic tasks such as:
  - *Study Methods and Comprehension*
    - Most students (**71%**) use AI to summarize assigned readings and study materials, highlighting how common it is used as a study tool.
  - *Support for Prep and Research*
    - Almost half of those using AI to summarize use these tools to condense research papers or lecture notes (**40-45%**).
  - *Less focus on Writing Revisions*
    - Only **37%** use AI to summarize their own writing, indicating it's applied more for comprehension than personal editing.

AI Summary Benefits: Time Saving

- Data reflects responses from **871** students who reported using AI for summarization
- *Productivity and Time Saving*
  - Nearly eight in ten students said AI helps them save time **(78%),** and **over half (52%)** said it helps identify key information
- *Feedback and Workflow Support*
  - **41%** use AI to organize their work, **35%** said it helps avoid busywork, and **30%** have AI produce feedback.

AI Helps to Edit Code

- **426** students (**33%**) who reported using AI for coding and programming assistance were asked how they apply it across different academic tasks.
- *Providing Feedback and Editing*
  - The majority of students use AI **(65%**) to diagnose bugs in their code and brainstorm **(46%)** alternative code methods
- *Drafting and Completing Code*
  - **37%** of these students utilize AI to produce their first coding drafts, and only **29%** use AI to produce full code outputs.

AI Coding/Programming: Saves Time and Avoids Mistakes

- *Time Saving*
  - **66%** of students who use AI to code report that it saves them time and **25%** said it helps them avoid busywork
- *Feedback and Iteration*

    - AI is most often used by students to prevent mistakes **(53%)**, followed by providing coding feedback **(47%)** and generating novel code or ideas **(45%)**

Key Motivations Behind AI Use

- *AI Saves Students Time*
    - Over **two-thirds** of students use AI in writing, summarization, and coding because it saves time. Specifically, **66%** of those who use AI for writing, **77%** for summarization, and **66%** for coding identify time efficiency as a primary motivation.
- *Provides Feedback and Support*
    - Receiving feedback was viewed as a key benefit among **47%** of the students who use AI to code, **30%** who use AI for summarization, and **72%** who use AI for writing.
- *Completing an Entire Task*
    - In contrast, a smaller share of students use AI to fully generate outputs. **28%** use AI to produce a first written draft, and **14%** to create entire texts.
    - In coding, **37%** use AI to generate initial drafts, while **29%** rely on it to complete full code outputs.

Perception Gaps: Students Overestimate Peer AI Use and Think Their Peers Violate Policy

One of the most noteworthy findings about student perception is that they believe their own AI use and likelihood of policy violation is far lower than that of their peers. While only **20%** of students say they always or usually use AI, **74%** believe their peers do use AI that frequently. The majority of students (**61%**) use AI sometimes or rarely, highlighting most engagement is selective rather than constant. Similarly, students not only overestimate AI use but also underestimate how often their peers follow AI classroom policies: **78%** of students report always or usually adhering to AI policies, while only **6%** of participants rarely or never follow AI policies. In contrast, **80%** of respondents perceive their peers as **never, rarely, or sometimes** following AI policies. Only **2%** of students believe other students always follow AI policies.

Although it is conceivable these results reflect a social desirability response (individuals underreport practices seen as negative), it is equally possible that the phenomenon is the same that leads students to believe their peers drink more alcohol than they do.[5]

Figures IV.2 and IV.3 below display the results about individual vs. peer AI use and individual vs. peer policy violations, respectively.

[5] Perkins, H. Wesley, Michael P. Haines, and Richard Rice. "Misperceiving the College Drinking Norm and Related Problems: A Nationwide Study of Exposure to Prevention Information, Perceived Norms and Student Alcohol Misuse." *Journal of Studies on Alcohol* 66, no. 4 (2005): 470–79. https://go.gale.com/ps/i.do?p=AONE&sw=w&issn=0096882X&v=2.1&it=r&id=GALE%7CA136122708&sid=googleScholar&linkaccess=abs.

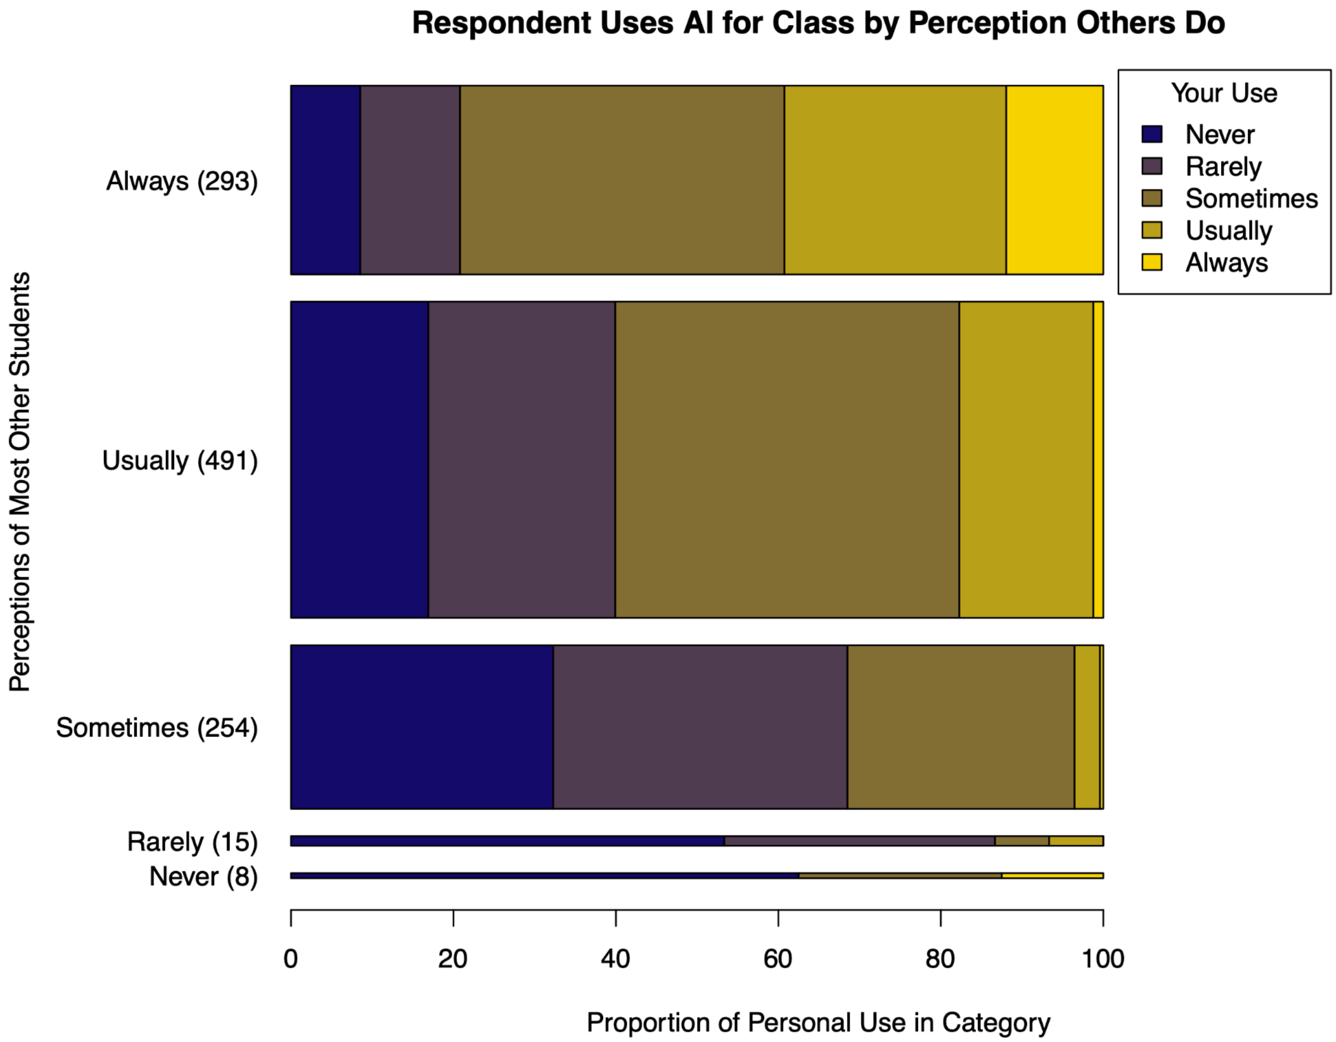

Figure IV.2 shows how much students report *using AI for their classes* (x-axis) contrasted with their perception about how much other students do (y-axis).

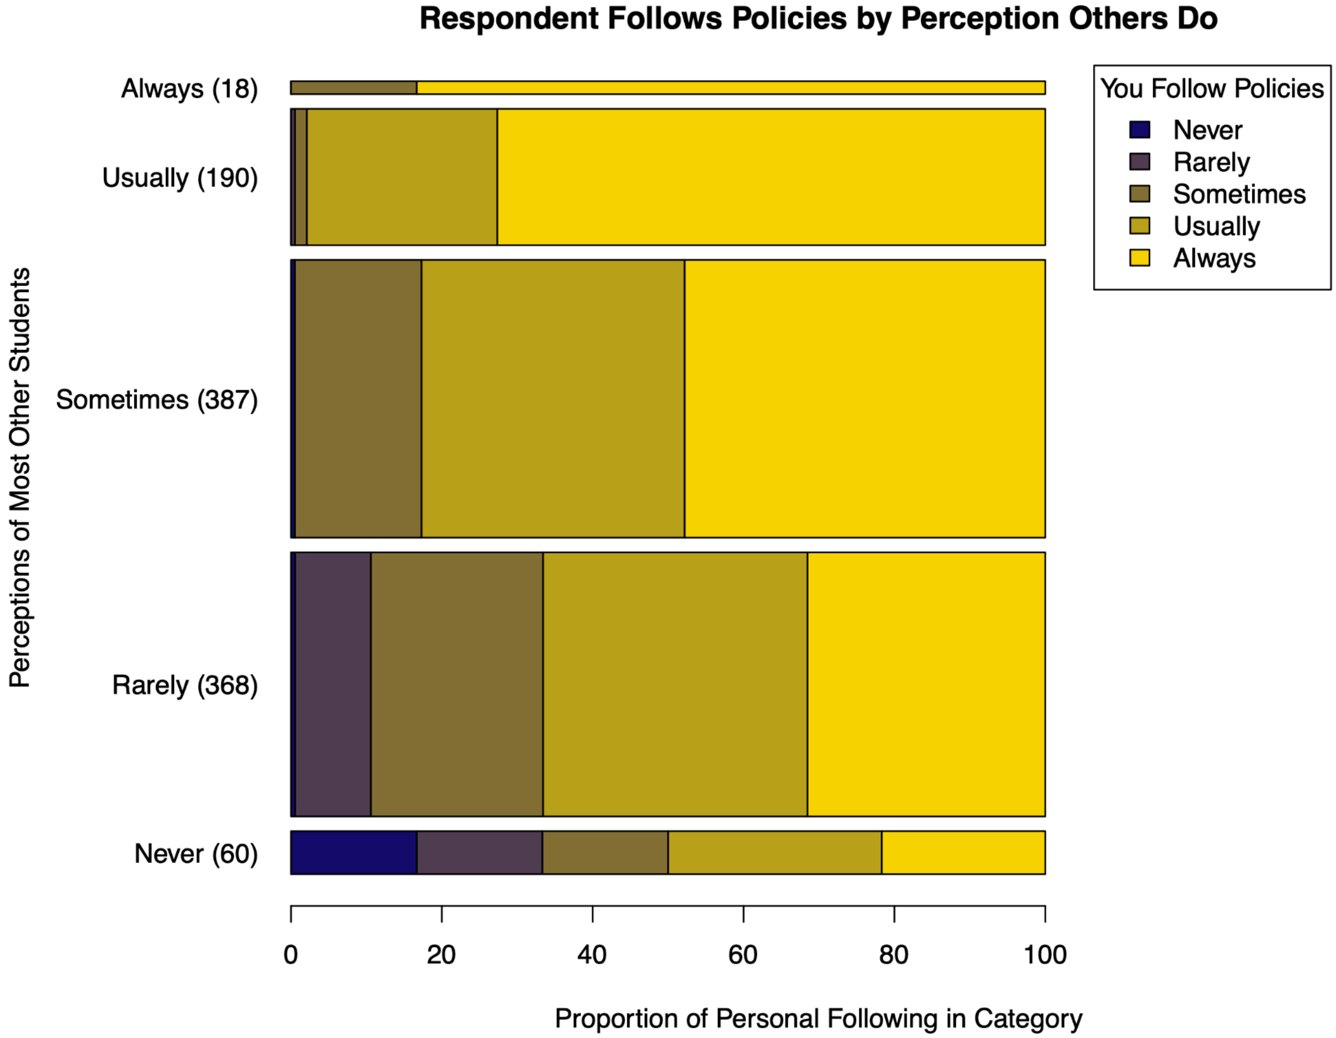

Figure IV.3 shows how much students report *abiding by AI policies* (x-axis) contrasted with their perception about how much other students do (y-axis).

Benchmarking Study

The working group reviewed the landscape of AI in education outside the University of Michigan. For this review, it commissioned Dr. Frank Hu, a data scientist at the Michigan Institute for Data and AI in Society (MIDAS).

The study covers 72 research universities and, in a separate analysis, a group of eight engineering-specialized institutions. The latter were included despite their significant differences from U-M (highly specialized, mainly very small) to get a sense of whether technologically oriented institutions approach AI differently. Data comes exclusively from publicly available websites.

Both analyses suggest the following:

- **All institutions are developing responses to AI in education** (as distinct from AI for research or administrative uses), but there is significant variation on the adoption of tools, policies, coursework, and requirements. The technical orientation of an institution does not predict any of these elements.
- **Mandatory AI requirements are rare**. Only 4 of 72 institutions require AI literacy for all undergraduates (Ohio State, Purdue, Stony Brook, SUNY Buffalo). Ohio State is the only institution that pairs a universal mandate with a named, funded provost initiative and published learning outcomes.
- **Competency frameworks are polarized and the most variable dimension**. Exactly 37 institutions have a CTL-level AI competency framework, 33 have none, and only 2 (Ohio State, University of Florida) have an institutional-level framework linked to graduation policy.
- **Formal assessment is essentially absent**. University of Florida is the only institution with a fully deployed university-level AI competency assessment. Seven others have announced instruments in development. The remaining 64 institutions assess AI learning at the instructor level only or not at all.
- **AI tool provision is near-universal, but FERPA compliance documentation remains a dividing line.** Nearly all institutions (69 of 72) provide students and instructors with at least one institutional AI tool — typically through enterprise agreements with Microsoft, Google, or OpenAI. However, only 40 have documented enterprise-grade FERPA compliance confirmed through institution-authored policy language. The remaining 29 provide tools without explicit published FERPA documentation. Three institutions (University of Washington, UC Santa Barbara, UC Santa Cruz) have no confirmed institutional AI tool provision as of early 2026.

The full analyses, including all data that informed them, are available in Google Drive for U-M Community members.

It is important to note that the study is a snapshot of one moment in a fast-developing environment. As the study shows, most universities are addressing AI in education in various ways. There are, nevertheless, worthwhile lessons for understanding how other universities have approached AI and where the U-M fits in that landscape.

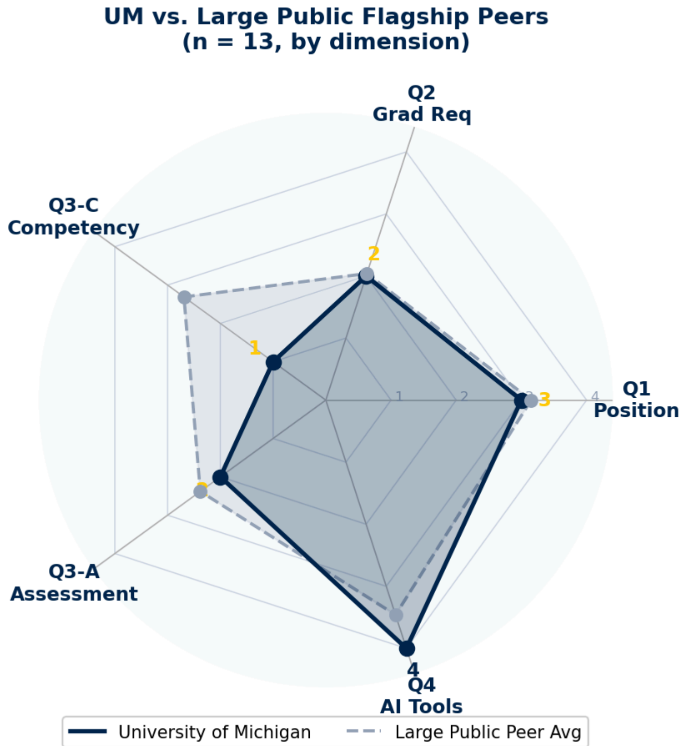

Figure IV.4 shows how U-M compares to other large public peer universities in terms of its publicly articulated position on AI (Q1), the existence of AI graduation requirement (Q2), articulation of AI competencies (Q3-C), assessment of competencies (Q3-A), and the the provision of AI tools (Q4). A higher number, i.e., being closer to the outer edge of the pentagon, indicates a more AI-prepared position.

V.A Teaching and Learning with AI

This section is about AI as a tool in teaching and learning — how instructors and students use AI to pursue learning outcomes in courses across the university. It is not about learning about AI; that is the subject of the AI Literacies and Competencies section. A political science seminar that uses an AI model as a debate opponent is in scope here; an Intro to AI course is in scope only for its pedagogy, not its subject matter.

The landscape of potential uses is large and growing quickly. We begin with a framework for evaluating any such use, then apply the evaluation framework to a catalog of potential uses.

1. An Evaluation Framework: Benefits; Harms; Design Choices

The overall evaluation criterion is simple: effective AI use by a student is use that improves the student's learning; and effective AI use by an instructor is use that improves students' learning. That criterion, applied to any particular proposed use, can be made more tractable with three questions:

- **Potential benefits:** How might this use help learning? *(Scaffolding? Practice density? Faster feedback? Access to expertise? Motivation?)*
- **Potential harms:** How might this use harm learning? *(Illusion of competence? Productive struggle short-circuited? Erosion of student agency or confidence? Overfitting to AI conventions?)*
- **Design choices: Are there ways to amplify the potential benefits and diminish the potential harms?** *(Sequencing in the course? Constraints on the AI's role? Reflection prompts?)*

The evaluation framework is discipline-agnostic by design; what it produces will look different in an intro programming course than in a writing-intensive humanities seminar, because learning goals, professional norms, and epistemic practices differ. At Michigan — where there is no central AI teaching policy, and, on this view, should not be — application happens in departments and courses, by the people who know the learning goals, the students, and the field. It is a decision tool, not a mandate to use AI.

The framework is equally a lens students can apply to their own AI use. The harms question speaks directly to a worry that surveys suggest many students already share: that reliance on AI will erode their own sharpness. Framing this concern as one of three co-equal questions, alongside potential benefits and design choices, keeps it legitimate without letting it dominate the evaluation. In practice, a student applying the framework to their own use might ask: Does this help me think more deeply, or help me avoid thinking? Am I learning the concept, or just completing the task? Can I reproduce or explain this independently afterward?

Whether the user is an instructor or a student, the framework deliberately excludes several common justifications for AI use. Efficiency, convenience, novelty, vendor features, and the observation that "students are already using AI" do not by themselves establish that a use improves learning. Efficiency improvements may be an indirect pathway to improved learning: freed instructor or student time might be used for other activities that enhance student learning. But there is no guarantee that saved time will be applied in that way. We focus here on direct impacts on student learning.

2. A Catalog of AI Roles

The catalog below distills the many AI uses being practiced and proposed in education, organized by the pedagogical function AI plays rather than by the underlying tool or discipline. It is not a list of endorsements. We analyze each role using the previous section's framework questions: potential benefits, potential harms, and design options that could amplify benefits and reduce harms.

- **AI as explainer.** Student asks a question; AI produces an explanation.
  - *Potential Benefits:*
    - lowers the cost of asking basic questions;
    - on-demand explanation at the student's pace;
    - access to analogies or perspectives beyond a single instructor's repertoire;
    - students embarrassed to ask in class can ask freely.
  - *Potential Harms:*
    - passive consumption that substitutes for the student generating the explanation themselves;
    - hallucinated explanations that read as authoritative;
    - illusion of competence, that having read an explanation equals understanding;
    - uncalibrated to course-specific terminology or examples.
  - *Design:*
    - requiring students to paraphrase or apply the explanation before moving on;
    - AI instructed to explain at multiple levels or via multiple analogies;
    - course-specific knowledge bases;
    - instructor visibility into which questions recur.
- **AI as Socratic questioner.** AI asks the student questions and probes their answers; the student does the articulation work.
  - *Potential Benefits:*
    - forces retrieval and self-explanation (both well-supported in learning science);
    - unlimited patience for repeated elicitation;
    - low-stakes space for students to practice articulating their thinking.
  - *Potential Harms:*
    - questions pitched at the wrong level or off-target;
    - AI may accept wrong answers or fail to identify confusion;
    - exchanges can feel like a game rather than learning;
    - frustration when the line of questioning doesn't match the student's need.
  - *Design:*
    - course-specific learning objectives supplied to the AI so its questions align;
    - requiring students to write a summary of what they concluded after the session;
    - instructor review of transcripts to catch miscalibration;
    - combining with instructor-led discussion so students see expert questioning modeled.
- **AI as a thought partner.** The user (instructor or student) brings a nascent idea, argument, problem, or solution approach. This extends the familiar idea of a *sounding board* — a passive listener that encourages you to express ideas and reflects them back to you — but goes further: the AI can actively point out weaknesses and connections, and suggest reformulations and extensions. This is work that happens before there is any artifact concrete enough for a critic to evaluate.
  - *Potential Benefits:*
    - articulation and externalization sharpen nascent ideas;

        - the AI's reformulation often surfaces structure the user had not yet found;
        - always-available iteration at the fuzzy front end of work, before anything is concrete enough for a critic;
        - holds a complex argument across long conversations without tiring.
    - *Potential Harms:*
        - **Ineffective partnership**:
            - sycophancy — the AI reinforces the user's direction rather than resisting, leading to premature convergence;
            - blind acceptance — the student accepts the AI's suggestions without a critical eye, again leading to premature convergence;
            - cliché import — the AI pulls in off-the-shelf framings, muddying rather than clarifying the student's ideas.
        - **Long-term atrophy**: capacities that only emerge from sustained solo practice and productive struggle may never form, including the ability to think alone when partnership is unavailable and developing a personal authorial voice.
    - *Design:*
        - instructing the AI to push back and disagree explicitly, not just reformulate;
        - requiring the user to restate the AI's reformulation in their own words before accepting it;
        - coaching students to engage in productive struggle *with* the AI thought partner.
        - alternating with unassisted work to check whether the ability to articulate ideas survives without the partner.
- **AI as practice-material generator.** AI produces practice items — problems, worked examples, case scenarios, synthetic data sets — for students to work on.
    - *Potential Benefits:*
        - unlimited variety at near-zero marginal cost;
        - items targeted to specific learning objectives or difficulty levels;
        - can feed into adaptive spaced-repetition systems, delivering classic learning-science gains (retrieval, spacing) at scale.
    - *Potential Harms:*
        - items that miss the distributional realism of real practice — edge cases, messy features, real-world variation;
        - encoded biases in default framings, demographics, or scenarios;
        - students calibrating to AI-item patterns that don't match what real work looks like.
    - *Design:*
        - instructor review and curation before release;
        - mixing AI-generated items with instructor-generated items;
        - student critique of generated items for realism, bias, and coverage gaps;
        - students iteratively guide the AI to improve items;
- **AI as producer of artifacts to critique.** AI produces an artifact — a draft, solution, argument, worked example; the student critiques, evaluates, or improves it.
    - *Potential Benefits:*
        - editing and evaluation practice, often higher-order than generation itself;
        - cultivates critical judgment and domain expertise;
        - exposes students to multiple plausible approaches via AI-generated variants;
        - inverts the usual "student produces, instructor evaluates" flow so students practice evaluation directly.
    - *Potential Harms:*

        - rubber-stamping or lightly editing AI output rather than engaging critically;
        - subtle AI errors that students miss, reinforcing wrong patterns;
        - critique reduced to template-matching;
        - if the AI output is already near-correct, real evaluation work disappears.
    - *Design:*
        - require student to articulate evaluation criteria;
        - deliberate error injection in AI outputs;
        - calibrating AI output quality to force genuine judgment;
        - requiring students to justify which edits to make and why;
        - assessing student judgment on AI output, not just the final revised product;
        - requiring identification of what is good alongside what is weak.
- **AI as critic.** Student produces work; AI evaluates, comments on, or gives feedback on it.
    - *Potential Benefits:*
        - always-available feedback — mid-revision, off-hours, without waiting on office hours or TA cycles, and therefore especially valuable where human feedback is bottlenecked;
        - lower social cost of being corrected by an AI than by a peer or instructor, so students who hesitate to expose incomplete work seek feedback more freely.
    - *Potential Harms:*
        - **Bad feedback**
            - uncalibrated to the course or task;
            - the AI's own taste or quality judgments misaligned with what actually counts as good work in the field or genre;
            - rubric-overfitting that loses sight of the work's larger trajectory;
            - mistaking volume of comments for depth of feedback.
        - **No internalization of feedback**
            - students ignore feedback;
            - students act on specific fixes without drawing larger lessons that transfer to future work.
    - *Design:*
        - layered human+AI feedback (AI for early drafts, human for judgment calls);
        - AI feedback scaffolded by a course-specific rubric;
        - requiring students to write a response to AI feedback explaining which suggestions they accepted or rejected and why;
        - instructor spot-checks for systemic miscalibration.
- **AI as simulator.** AI constructs a bounded scenario for the student to engage with — a character to converse with (mock client, mock patient, historical interlocutor, debate opponent) or a world to explore (a historical setting, an ecosystem, a market, a legal system). The student acts within the scenario; the AI responds in character or in-world.
    - *Potential Benefits:*
        - cheap, repeatable access to semi-authentic practice contexts;
        - low-stakes practice on high-stakes scenarios.
        - access to settings that are inaccessible in reality (dangerous, historical, hypothetical, or counterfactual.)
    - *Potential Harms:*
        - succeeding in the simulation without developing the skill needed for real situations;
        - simulations may miss dimensions the real thing has — affective and social in character role-play, material and physical in simulated worlds;

        - verisimilitude may disguise simplifications the student should recognize.
    - *Design:*
        - grounding scenarios in real cases, situations, or systems from the field;
        - instructor review of transcripts or session traces;
        - pairing simulations with live-actor, real-world, or empirical observation;
        - reflection on what the AI got right and wrong about what it was simulating.

3. Recommendations

These recommendations primarily support courses that integrate AI in some way. For courses that choose not to, the main support needed is around assessment redesign. That is addressed in the Assessment and Academic Integrity section.

- **[For instructors]** Bring the framework into the course. Faculty adopt one of two approaches:
    - A course activity in which students apply the framework questions (potential benefits and harms, and design choices that amplify benefits and minimize harms) to AI uses relevant to the course.
    - An AI posture statement in the syllabus — not an "AI policy" listing permitted and forbidden uses, but an articulation of the instructor's design choices, the rationale behind them, and how those choices connect to the role AI plays in the discipline. Writing such statements forces instructors to reason through their choices rather than inherit defaults, and gives students an explicit account of the pedagogical logic of the course.
- **[For the university]** Make it easy to try innovations in teaching and learning
    - Work out logistics of student API access
        - Lab fee? University pay? Data privacy issues.
    - Sandbox for innovation in Canvas
        - Easy to do a pilot install of a third-party tool in Canvas for just your section
        - Easy for instructors to get agents to act on their behalf in Canvas
- **[For the university]** Support corps of undergraduate Learning Experience Designers (LXDs) who help instructors. Many of the §2 patterns are difficult for an individual instructor to implement well without assistance. Partnering with an AI-fluent LXD puts the help inside the course design process rather than requiring every instructor to develop fluency from scratch.
- **[For the university]** A knowledge base and community of practice for sharing innovations. LXDs and instruction innovators will benefit from learning from each other. The University of Michigan should evaluate whether to build its own knowledge base or contribute to a cross-institutional effort. Either way, staff support would be needed to convene and sustain the community of practice.

V.B Assessment and Academic Integrity

Academic Integrity and Purposes of Grading

Academic integrity is central to the mission of the University of Michigan, to the value of a U-M degree and to the development of a robust intellectual community among faculty, staff and students. It enables trust within the academic community, supports a positive learning environment, and ensures that academic work reflects individual knowledge and effort.

For students, maintaining academic integrity supports the development of skills and proficiency expected by employers, graduate programs, and professional communities. It reflects respect for peers and instructors by upholding fair standards. It also supports learning through meaningful engagement with course material. It further reflects maturity and sound judgment in navigating ethical decisions beyond the classroom and fosters pride and confidence in one’s own abilities.

When academic integrity breaks down, trust erodes and fair evaluation becomes more difficult. It creates inequity by giving some students an unfair advantage while discouraging others. A learning culture without academic integrity bypasses the development of key skills. Over time, this weakens the meaning of grades and the value of the credential itself.

These considerations directly inform the primary purposes of grading: to support learning, communicate competency, and evaluate student improvement. These purposes are grounded in clearly defined course objectives and learning outcomes. Effective course design begins by identifying what students should learn and be able to do, and aligning assessments and grading practices with those goals. This backward design approach ensures that grades reflect progress toward meaningful learning outcomes.

At the University of Michigan, we seek to build a culture of academic integrity among faculty, staff and students.

1. Challenges in Evaluating Student Learning

The integrity challenge is real, qualitatively new, and not solvable through detection or assignment design alone.

The majority of faculty report difficulty in accurately assessing student knowledge. While AI creates new opportunities for rethinking assessment design and grading, it also introduces structural challenges for evaluating specific learning objectives and skill development. These challenges put stress on the purposes of grading outlined above. The value of grades depends on the shared assumption that assessment evidence reflects what a student can actually do. In many unproctored contexts, that assumption no longer holds.

Historically, academic misconduct has existed in the form of paper mills, plagiarism, and impersonation. Detection was imperfect and enforcement uneven, but access to these options was limited and the quality of the fraudulent work varied widely.

The current environment is qualitatively different. High-quality artifact production is now free or low-cost, immediate, and embedded in everyday tools. AI is readily available and often integrated directly into platforms students already use (e.g., U-M GPT, AI assistants in browsers and library ebooks). The decision to use it does not require deliberate planning or effort; it is often the default next step. Faculty also report a loss of control over assessment, as practices that once reliably measured student learning no longer do so.[6] Responsibility for addressing this shift is often placed on individual instructors without corresponding institutional support or guidance.

In this context, traditional unproctored assignments such as essays, problem sets, lab reports, and take-home exams no longer reliably reflect student mastery or produce the engagement they were designed to support. As a result, both the incentive function of grading and its role as a signal of competency are weakened.

Assessment of student competency requires evidence that responses are the student's own. Proctored environments, whether in person or through carefully designed alternatives, provide direct evidence of student performance. Proctored environments support trust in the evaluation process, including students' confidence that results are fair and reflect their own work. At the same time, the use of proctoring raises important considerations related to access, equity, and student experience, particularly for remote or online learners.

A range of alternative approaches has been proposed to address these challenges. While some offer value, none provides a stable or scalable solution on its own.

AI detection, both human and machine, is limited. In practice, error rates are too high for consequential decisions, with false positives disproportionately affecting non-native English writers and students with atypical styles. As LLMs improve, the distinction between AI-generated and human-produced work continues to narrow. Users of LLMs also become better at avoiding detection sensors, limiting the long-term viability of detection as a foundation for academic integrity.

Lockdown browsers and remote proctoring aim to control the student's device during assessment. Lower-intrusion approaches restrict on-screen activity but do not account for off-device resources, such as phones or secondary devices. Higher-intrusion approaches introduce camera monitoring, audio recording, and behavioral tracking. While these may approximate proctored conditions, they raise concerns related to privacy, equity, and demographic bias in AI-based behavior monitoring. They also undermine student trust in the assessment process, weakening the community ethic of academic integrity. These approaches introduce substantial operational costs, including licensing, per-exam fees, and administrative overhead. For these reasons, they are not a universal solution. Their use may be appropriate in limited contexts where in-person alternatives are not available (e.g., fully online degrees) and the level of intrusion is proportionate to the assessment. Separately, device-control tools can serve as one layer within a proctored classroom, such as restricting use of additional software.

Process-trace analysis focuses on how work is produced, using keystroke patterns, revision histories, and time-on-task data to infer authorship. However, AI systems can already generate plausible writing traces, and as these tools improve, human and AI-generated processes are likely to become increasingly indistinguishable. These approaches may offer short-term benefits but do not provide a durable solution.

Echo-checks, or proctored follow-ups on unproctored work, assess whether a student can explain or defend a submitted artifact. While they can confirm comprehension and some learning processes, they do not establish authorship. A student who studies an AI-generated response may perform as well as one who produced the work independently. As a result, echo-checks function more effectively as standalone assessments of comprehension rather than as verification of ability to perform a particular task.

---

[6] The reliability of conventional instruments to assess *actual* learning has always been a difficult question, but the situation introduced through AI is different because it can be impossible to tell whether the student even did the *work* the assessment asked for.

Taken together, these limitations point to a broader conclusion. The challenge is not primarily one of detection or assignment design, but a structural shift in what unproctored work represents. Unproctored artifacts no longer reliably provide evidence of individual competency. Assessment practices that continue to treat them as such risk measuring something other than what they intend to measure. Addressing this shift requires reconsidering where and how assessment evidence is generated, as well as clarifying what learning outcomes are being measured and aligning assessment practices accordingly.

2. Rethinking Traditional Assessments

Proctored assessment is a design space, not a format. Large-enrollment infrastructure should expand access to that space.

Readers often associate “proctored assessment” with synchronous, 60-120 minute midterm and final exams composed of multiple-choice questions (MCQ), short-answer items or short essays. It is problematic to rely on such exams as the only form of assessment. They concentrate grade weight into a few high-stakes events, raising per-event anxiety and making grades brittle to single-day performance. And they reach only a narrow band of skills, those demonstrable in a timed individual sitting.

Fortunately, the traditional midterm or final exam is just one point in a much larger design space, which can be described along **four independent dimensions**:

- **Item types** include, for example, MCQs, short answers, essays, code, performance-based tasks.
- **Time constraints** vary from fixed to flexible. In fixed settings, a shared time block applies to all students, with formal accommodations (e.g., 1.5×, 2×) handled as formal exceptions. In flexible settings, students work at their own pace within defined deadlines.
- **Personalization** refers to how prompts vary across students. At low levels, all students receive the same items or minor variations (e.g., randomized from a shared pool). At higher levels, items are generated dynamically or tailored to the student’s prior work.
- **Interactivity** refers to how later questions depend on earlier responses. At low levels, assessments are static and one-shot. At higher levels, they include adaptive follow-up. These may include oral exams.

Most courses rely on conventional exam formats because other assessments *(e.g., tutorials, oral exams)* are difficult to scale, requiring more instructor time and attention. These approaches have long been used in small-enrollment settings, but extending them to large courses introduces practical constraints.

Each dimension places demands that scale poorly in conventional classroom settings. Rich **item types** often require specialized environments (e.g., oral exam, simulation-based, coding in controlled environments, markup of supplied materials) and can increase grading demand. Flexible **time** structures are difficult to support within fixed, shared schedules, and classroom access. **Personalization** requires systems that can deliver different items to different students and maintain sufficient pools or generate tailored prompts around individual student input. **Interactivity** requires real-time evaluation of prior responses and generation of tailored follow-up questions, which a single instructor cannot provide at scale.

Four infrastructure investments can address these challenges.

- **Assessment centers.** Dedicated, scheduled environments where students complete proctored work on instructor-authorized devices. These settings support a wider range of **item types**, including aurally isolated rooms for oral formats, controlled computer environments with authorized-tool access for simulation-based tasks, dataset analysis, coding tasks, and markup of supplied materials. They also enable flexible **time structures** through individual-pace scheduling decoupled from shared class meetings, identity verification, any necessary accommodations, and retake logistics provided as standard features rather than negotiated per instructor and per student. Digital delivery systems support personalization by routing different items to different students. When combined with non-synchronized scheduling, this approach also reduces integrity risks associated with students sharing information across sessions.

- **Tutorial-section infrastructure**. Weekly tutorials in which students present and discuss their work with a GSI in small groups are a recurring proctored format. They do not need the identity verification and controlled computer environments that a testing center provides. However, their widespread adoption will require registrar workflows and scheduling small conference rooms as classrooms. The STATS 413 tutorial pilot (see vignette) navigated this with ad-hoc instructor effort.
- **AI in the proctored-assessment infrastructure.** AI can function as infrastructure rather than student shortcut, enabling capabilities that human-only processes cannot deliver at scale. Two main roles are relevant, each with distinct benefits, risks, and design considerations.
    - *AI-assisted grading of open-ended responses.* This supports rich **item types** at scale and enables **interactivity** through real-time evaluation of student responses needed to trigger appropriate follow-ups.
        - Potential benefits include enabling non-MCQ formats at large courses, consistent rubric application across large numbers of responses, rapid turnaround that supports retake and mastery-oriented designs, and real-time grading of in-session responses to support adaptive follow-up.[7]
        - Risks include calibration drift over time, demographic bias in grading, students learning to optimize for the grader rather than the rubric, overreliance on surface-level rubric compliance, erosion of human oversight at scale, and negative student perceptions which erode a community ethic of academic integrity.
        - Effective design requires safeguards such as human sampling audits on a fixed fraction of responses, student-facing appeal processes, calibration checks against human-graded subsets, full transparency about which components are AI-graded, bounded use in contexts where the rubric is genuinely operationalizable.
    - *AI-assisted item generation.* This supports **personalization** by enabling large pools of item variants and on-demand item generation tailored to individual student input. When combined with real-time grading, it also supports **interactivity** through adaptive follow-up probes keyed to the student's prior responses.
        - Potential benefits include rapid generation of many variants aligned with specified learning objectives, supports **personalization** at scales infeasible with human-only item writing, tailored probing based on individual student input, scales viva-style adaptive questioning beyond what individual human examiners can deliver.
        - Potential harms include hallucinated or inaccurate content, items that look plausible but miss the intended learning target, embedded biases in framings or scenarios, and drift toward AI-generated conventions rather than authentic disciplinary practice.

[7] For example, experiment using a randomized control test design suggests that AI-mediated student feedback can improve student performance. See Lu, Xinyi, Kexin Phyllis Ju, Mitchell Dudley, Larissa Sano, and Xu Wang. "AI-Mediated Feedback Improves Student Revisions: A Randomized Trial with FeedbackWriter in a Large Undergraduate Course." arXiv:2602.16820. Preprint, arXiv, February 24, 2026. https://doi.org/10.48550/arXiv.2602.16820.

        - Design to mitigate these harms requires instructor review and calibration of sample items before deployment, validation of cross-variant equivalence, red-teaming for bias and content error, and bounded use in which AI-generated items are reviewed and approved by instructors.
- *Support corps of undergraduate Learning Experience Designers (LXDs)*. Instructors require support to design assessments in richer regions of the design space. The same LXDs who help instructors incorporate new AI-assisted pedagogies (see the Teaching and Learning section) can also help them design assessments.

3. Accreditation, Compliance and Reputation Considerations

Developing institutional infrastructure for assessment addresses an institutional-risk question, not only a pedagogical one.

Outcomes assessment. Regional and programmatic accreditors *(HLC, ABET, AACSB, LCME, and others)* require evidence that students attain stated learning outcomes. Evidence has historically included coursework and assessment. If assignments are no longer reliable evidence of student-produced competency, programs will need proctored components in their outcomes-assessment plans. Schools and departments should begin to consider realigning learning outcomes and assessment modalities.

Credential meaning as institutional risk. A University of Michigan degree signals “Leaders and Best.” This depends on external parties — employers, graduate programs, licensing boards, ranking agencies — trusting that grades track competency at an elite institution. Reputation harm is slow-moving and hard to reverse.

Federal Title IV / distance-education context. Identity verification in distance education is already a federal requirement. LLMs raise the stakes of the meaning of “the student did the work” for online and hybrid programs in particular; regulatory expectations are likely to tighten, not relax.

First-mover dynamics. Institutions investing early in proctored infrastructure *(assessment centers, tutorial-style teaching support, faculty development for proctored formats, hiring of proctors, graders and additional SSD staff)* will make stronger claims to accreditors, employers, ranking institutions and prospective students. Laggards face a scramble when external pressure arrives, with less time to design well.

4. Summary

Academic Integrity is a valued community ethic at the University of Michigan that ensures a culture of “Leaders and Best.” The integrity challenge around assessment is a novel challenge in the age of AI and LLMs. Students need to be able to trust that their work is being graded fairly, relative to their peers. Students are concerned about false positives in detection systems and they lack trust that their peers are completing their work honestly. Instructors need to be able to assess student competency and skills accurately. The university should guard its reputation by guaranteeing the authenticity of a University of Michigan degree. Competency certification must rest on proctored activity — no output-based detection tool, remote proctoring setup, process-trace scheme, or Echo-Check on unproctored work will fully preserve integrity in the long run.

Proctored assessment is a design space, not a single format. Large classes have depended on conventional-exams because some assessments have been impractical at scale.

Four infrastructure investments — assessment-center capacity, support for tutorial sections, AI in the proctored-assessment infrastructure, and a support corps of learning experience designers — would open much of the design space to broader use. These investments would be non-trivial institutional undertakings across all three campuses. The time is now to start pilot tests and planning for potential broader implementation.

Vignette: Oral Assessment in Statistics

In STATS 413 (fall 2025), SYK piloted an alternative to traditional labs that addresses the problem of AI one-shotting homework. Instead of labs, the course ran weekly tutorials: groups of five students meeting for 50 minutes with a GSI. The primary activity was students presenting their solutions to the previous week's homework to their peers. The motivation for this format is it frees the students to use AI to help them come up with the solutions as long as they internalize the help from the AI. The GSI grades the presentation, ensures every student gets a chance to present, and asks clarifying and follow-up questions designed to probe genuine comprehension.

Key logistics:

- Tutorials counted for 20% of the overall grade. The problem sets themselves were not graded.
- Each GSI ran 8 tutorials per week (8 × 5 = 40 students), preserving the standard 1:40 GSI-to-student ratio.
- Tutorial supervision replaced the homework-grading portion of GSI work; weekly hours remained well under 20 hrs/wk.
- Students were assigned to tutorial time slots via a centralized optimization that minimized total dissatisfaction across student preferences (see How we assign students to tutorials | STATS 413). It would have been preferable to let students sign up for tutorial time slots through the registration process but that proved to be logistically difficult for the registrar.
- Tutorials were held in the lab classroom and in the statistics office-hours classroom; ideally a dedicated room would be reserved for the full tutorial day.

Midterm evaluations (75/164 students responded; 46% response rate) were generally positive. The most common student suggestions:

- More time per problem — focus on one homework problem per tutorial rather than covering all of them (5/35 respondents who gave suggestions).
- Some review of course material alongside presentation (5/35).
- Concern that the graded format was anxiety-inducing and discouraged students from asking questions (4/35).

For instructors, once the format is established, tutorials require almost no additional preparation: the tutorial materials are the homeworks themselves.

V.C AI Literacy and Graduation Competency

This section is about AI as a subject of learning — what graduates should know, be able to do, and evaluate as AI users themselves. For how AI can be deployed as a tool within courses that aim at other learning outcomes, see the Teaching and Learning section above.

Establishing Baseline Expectations for U-M Graduates

U-M needs to prepare students to engage thoughtfully, creatively, and responsibly with AI systems as they become ubiquitous and deeply integrated into many academic, creative, and professional environments. While some students (e.g., STEM disciplines) may need to focus on technical mastery, all students should meet baseline expectations emphasizing critical understanding, creative practice, and ethical judgment. **All graduates are expected to move from passive consumption to actively shaping, directing, and evaluating AI systems — and, where they judge it warranted on that basis, criticizing them — in their roles as citizens, learners, and contributors.**

The three areas below each include ethical considerations rather than treating ethics as a separate fourth area. How AI systems are built and how human values shape their behavior is part of understanding; how AI is used appropriately is part of using; and judgments about whether and how AI should be deployed are part of evaluating.

1. Understanding AI (Conceptual Literacy and Critical Awareness)

Students must be able to comprehend what is happening "under the AI hood" and to think critically about AI design, utility, and limitations. They should:

- Understand how AI systems function as sociotechnical processes. This includes grasping the foundational systems pipeline, from data collection through model development to deployment, and the human choices at every stage that shape what the system does and whom it affects.
- Recognize that AI outputs are probabilistic predictions shaped by training data, not factual retrievals. For large language models in particular, recognize that the same prompt can produce different outputs across runs, and that the fluency of an output is not a reliable signal of its correctness.
- Be able to trace how human values, historical patterns, and cultural assumptions embedded in system designs and training data shape AI behavior, including the criteria for what counts as a good output, and errors and biases in those outputs.
- Develop "systems thinking" to understand how modern AI systems create emergent behaviors through relationships and feedback loops — and the questions of responsibility those behaviors raise when no single component fully owns the outcome.

2. Using AI (Creative and Collaborative Practice)

The practical application of AI focuses on shaping AI behavior and engaging it as a creative partner rather than just operating a generic tool. Students should:

- Learn to thoughtfully frame problems to determine when AI mediation is appropriate and useful.
- Be able to curate, contextualize, and critique data, recognizing it as a socially produced resource rather than neutral raw material.

- Be able to configure and adapt AI systems through experimentation, prompting, setting constraints, and designing workflows.
- Collaborate with AI iteratively, treating the interaction as an ongoing dialogue and refining systems based on unexpected outputs and failures.
- Integrate AI meaningfully into specific disciplinary workflows, adapting approaches to the specific questions, methods, and values of their field.
- Build the metacognitive skills necessary to enlist AI systems in acquiring domain-specific knowledge and identifying their own knowledge gaps.

3. Evaluating AI (Judgment and Accountability)

Graduates ought to be trusted as thoughtful, responsible human actors within AI-mediated systems. They should be able to:

- Critically evaluate whether AI should be used at all in any given context.
- Articulate and navigate complex ethical trade-offs, balancing competing values such as accuracy versus bias, efficiency versus equity, and capability versus environmental impact.
- Establish norms of honesty by disclosing AI involvement transparently in all academic and professional work.
- Recognize their own agency and accountability, understanding that technology amplifies rather than absolves human responsibility for decisions made with AI assistance.

V.D Ethics, Access, Equity

The faculty survey (Appendix F) shows that faculty views on AI are wide-ranging but much less polarized than public discourse might suggest. Faculty who may be called AI-skeptics or AI-adopters share some common concerns around equity, ethics and access. The university benefits by taking seriously critique and suggestions from all faculty, and looking at the most common concerns among diverse groups of faculty.

1. Responsible and Ethical Use of Generative AI

A supermajority of faculty in all disciplines and schools believe teaching students “ethical use of AI” is at least “somewhat important.” Over fifty percent of faculty in all disciplines believe it is “very important.” A supermajority of faculty in all disciplines believe it is “very important” to teach students about the negative impacts of AI. A supermajority of faculty believe that protecting student data is at least “somewhat important.” These findings underscore the importance of including ethical discussions of generative AI as a central part of AI teaching and learning.

- **Environmental cost.** Generative AI consumes meaningfully more energy and water than the computational baselines it is replacing. A course or program that scales up AI use is, in aggregate, externalizing an environmental cost. Institutional AI guidance documents should include an account of environmental trade-offs, and curricula that engage students as AI users should also engage them as AI consumers, with lifecycle costs made visible.
- **Bias.** A large number of AI systems have been shown to have racial, gender and other bias in output. For example, bias in facial recognition software has negatively affected communities of color. Detroit Police Department’s facial recognition program led to high false positives, wrongful arrests and a settlement that restricts its use. Meanwhile, Detroit Public Schools is now using facial recognition technology at school campuses, causing worry for community members. According to DPSS, U-M does not currently use facial recognition technology with campus security cameras. Students and faculty should be informed of any planned changes. Curricula should teach students about bias in AI systems. Students will not only be critically evaluating output, but some will also become designers of AI tools.
- **Data, labor, copyright and supply-chain transparency.** AI systems are built on training data, and on data-labeling and reinforcement-learning labor, whose provenance and conditions are often opaque. As an institution whose legitimacy rests on intellectual honesty about how knowledge is produced, U-M has both the authority to prefer vendors who are transparent about training data sources, data-labor practices, and model update cycles. Students should also understand the issues at stake around information provenance, copyright and fair use in changing legal and ethical contexts as it relates to generative AI.
- **Consent architecture for student and faculty work.** Student coursework and faculty intellectual output are, by default, subject to the terms of the platforms students and faculty use. Many commercial tools train on user-provided content unless users navigate opt-out settings. The institutional position should be that student and faculty work is *not* training material unless the author explicitly chooses otherwise, and that this protection is negotiated into vendor contracts. This also protects the intellectual property of students, faculty, and the institution. We also believe AI literacy education should teach students about consent and privacy.
- **Privacy, FERPA, and sensitive contexts.** AI use in courses involving student records, clinical data, or other protected information raises compliance questions that go beyond general ethics. These cases should be governed by clear institutional guidance rather than left to individual instructor judgment,

particularly in clinical, counseling, and research-methods courses where inadvertent disclosures are most likely.

- **Student cognition, mental health and wellbeing.** Data is only emerging about the psychological and cognitive effects of addictive technology and social media use. Very little is currently known about the possible side effects of young people's dependence on generative AI for learning or for social support. The university should support student wellbeing by including these discussions in courses focused on the ethics of AI, provide some space for screen-free learning, and should invest in research in the area of young people, wellbeing and technology.

2. Equity and Resource Disparities

According to the faculty survey, the majority of faculty in all fields find it at least "somewhat important" to mitigate pre-college AI preparation and access. A majority of faculty in all fields find it at least "somewhat important" to accommodate student disabilities using AI tools. A majority of faculty in all fields (except the humanities) find it at least "somewhat important" to make sure every student has access to AI tools.

The university licenses institutional services (U-M GPT, Maizey, the Google suite), and faculty can approve additional tools for use in their courses. Beyond this baseline, students vary widely in their ability to acquire premium subscriptions, their comfort with higher-capability tools, their knowledge about AI design, and their prior exposure to high quality generative AI learning or computer science training before arriving at U-M.

- **Prior-exposure gap.** Students arriving from well-resourced secondary schools, from AI-saturated tutoring environments, or from families in technical professions enter with substantially more hands-on AI experience than first-generation, rural, and international students. Policies that assume a common baseline disadvantage the latter groups, not in the superficial sense that they cannot use the tools, but that they may defer to AI outputs rather than interrogate them. Research on technology adoption and stereotype threat suggests that students from underrepresented groups are more likely to treat authoritative-sounding outputs as authoritative. Students may also feel like they can never compete with generative AI outputs and lack confidence in their own thinking. Building student epistemic confidence that they can challenge, override, or disregard an AI system is an equity intervention.
- **Language and cultural-register disparities.** Large language models are trained disproportionately on English-language text, and performance in other languages is measurably lower. Students writing in their second language, or in disciplinary traditions underrepresented in training data, may receive materially inferior assistance from the same tools. This has direct pedagogical implications for a course that assumes "AI is available to everyone equally."
- **Cost and account-risk disparities.** Many AI tools charge for their most capable versions, and price increases are being rapidly adopted. Courses that effectively require premium AI, whether explicitly or through assignment design, silently shift costs onto students, lock students into paying higher amounts in the future, and push them into creating third-party accounts whose data-handling practices the university has not vetted. The working group discussed how financial disparities are already appearing. When premium tools are required for courses, they should be explicitly included in course **requirements**, which makes them eligible for financial support. Students also need to know about the full suite of options currently available to them through the university.
- **Disability and access.** Access for disabled users is uneven. Some tools are made with disabled users in mind; creating greater accessibility in a wide variety of teaching and learning contexts. New tools should be tested with disabled users in mind and with faculty input. Meanwhile, some AI tools have not taken disabled users into account in the design process. Instructors need information on accessibility tools available at U-M and the limits of existing tools.

3. Clarity and Transparency for Students

Students need clear guidance and transparency for expectations around AI use.

- **Syllabus statements can help guide student use.** Faculty should include AI statements on syllabi so that students are aware of course-level expectations. Departments and faculty should discuss what kind of statements are useful in their disciplines. Shared language can help ensure students receive clear and consistent guidance across courses. Examples of syllabus statements can be developed by CRLT and other university centers.
- **Discussion of ethical AI use in published university academic integrity policies.** Colleges should develop discussions on the ethical use of AI in their academic integrity policies; colleges should also provide clear examples of academic dishonesty. ITS should develop guidance on privacy and safety for AI tools in their student-facing digital citizenship statement.
- **Non-punitive channels for uncertainty.** The working group heard repeatedly that AI-related integrity issues often originate in confusion rather than deliberate dishonesty. Students need a place to ask questions about compliance before submission, without that inquiry itself becoming adversarial. An accessible, consultation-oriented resource for AI-related uncertainty might be piloted.
- **Institutional transparency.** Universities that adopt AI policies without explaining their reasoning risk a credibility gap, particularly with students who are themselves sophisticated AI users and will quickly notice where policy diverges from institutional practice. The working group recommends that the university regularly publish a frank accounting addressed to students: progress on data centers; what AI tools are in institutional use and why; what data has been collected about student AI use and how it has been handled; how AI-related integrity cases have been adjudicated in aggregate; and what the institution has changed its position on since the prior cycle.

V.E Industry, Educational Technology, and Workforce Context

Generative AI is transforming how work is performed at the same time that it is transforming how students learn. The conditions under which students prepare for work, and the conditions under which they will eventually do that work, are converging — but asymmetrically. The AI capability frontier moves in months; vendor product cycles turn over every two years or so; undergraduate degrees run four years; institutional curricula and employer conventions evolve on cycles longer than any of these. Any educational posture calibrated to present entry-level expectations will be behind by the time a first-year student graduates, let alone mid-career.

This section addresses what it means for U-M to situate its educational strategy inside this convergence while preserving the university's distinct role. That role is not to mirror industry practice, nor to hold it at a distance. It is to prepare graduates whose judgment, domain competence, and intellectual autonomy remain useful on timescales that outlast any specific tool or workflow. The subsections that follow take up three aspects of that posture: what is actually shifting in entry-level expectations (§1); how AI should be positioned pedagogically as a system students supervise rather than operate (§2); and what conditions engagement with industry and vendors should satisfy to be defensible (§3).

The university's comparative advantage in an AI-mediated labor market is the cultivation, at a timescale employers cannot match once graduates are on the clock, of *the judgment that AI does not itself supply*.

1. Shifts in Entry-Level Expectations

Across sectors, tasks long considered foundational at entry level — drafting routine documents, assembling preliminary analyses, producing first-cut code, writing clinical documentation, summarizing meetings and research — are being absorbed into AI-augmented workflows. This has two implications that deserve to be distinguished, because they point in different directions.

- **The content of entry-level work is shifting.** Employers are placing greater weight, earlier in a career, on capabilities historically associated with mid-career professionals: framing problems rather than executing defined tasks; evaluating and refining the outputs of systems; exercising judgment about when a tool is and is not appropriate; coordinating across roles that now include both human and machine participants. The skills that differentiate graduates in this environment are less about operating AI tools — which the tools themselves are engineered to make easy — and more about domain competence, communication, and the ability to work with other people around AI-mediated outputs. Increasingly, the differentiating human contribution is the ability to understand collaborators, anticipate how they will respond, and synthesize information into conclusions on which other people can act. AI, in this view, is used to make the worker better at the task, not to replace the worker's own thinking.
- **The entry-level career ladder may itself be under pressure.** The routine production tasks being absorbed by AI were also the tasks through which entry-level professionals historically built tacit judgment. If those tasks disappear, so does the apprenticeship structure that produced mid-career competence. Whether organizations will develop new structures for building judgment in-house is an open question; in the interim, graduates may be expected to arrive with a readiness their predecessors developed over several years of routine work. This makes the university's role in building judgment *before* the first job more consequential, not less — and makes it risky for the institution to defer that cultivation to the employer.

- **At the same time, certain entry-level positions may incentivize blindly trusting AI outputs.** Pressures for throughput and efficiency in entry-level work may encourage rote AI usage and discourage critical thinking. Natural workplace competition and the need to present competence rather than willingness to learn risks neutralizing the university's emphasis on judgement over blind trust. To withstand these pressures, students must deeply believe in the *value* of human thought and practice-based learning and improvement. It is the university's responsibility to create learning environments that continue to value critical thinking and instill resilience against the path of least resistance.

These shifts create pressure on credential signaling that deserves to be named. When AI makes routine entry-level output cheap, the value of a University of Michigan degree depends more than ever on what the degree can credibly assert about the graduate's judgment, rather than on the graduate's capacity to produce artifacts. The assessment reforms in §V.B are, in this light, not only an academic-integrity matter but a labor-market one: the signal the degree carries is only as credible as the assessment practices that stand behind it.

Two cautions follow. First, educational priorities should not be reduced to what employers say they want this year. The half-life of specific employer asks — certifications in particular vendor stacks, fluency with particular tools — is shorter than the degree being awarded, and reacting to each cycle produces a curriculum perpetually behind. Second, "preparing students for an AI-mediated workforce" should not quietly become a justification for narrowing the curriculum toward whatever industry-adjacent skills are most legible in a recruiting brochure. The durable capacities graduates will need — judgment, communication, the capacity to work with and around other people — are as readily cultivated in a literature seminar as in a data-science course, and in some respects more so. The graduate-competency framework in §V.C reflects this view: it asks what graduates should be able to do across a career, not at the first interview.

2. AI as a Tool Students Must Evaluate and Supervise

Supervisory judgment about AI develops only by supervising AI in conditions where its failures matter.

The most consistent message from workplace practice is that AI functions less as a productivity tool in the calculator sense and more as a system requiring ongoing oversight. AI outputs are confident, plausible, internally consistent, and — not infrequently — wrong. They confabulate; they import biases from their training; they miss the details that matter most for a specific case; they fail in ways that are difficult to detect without domain expertise. Graduates who have not developed the habit of interrogating AI outputs will defer to them; graduates who have will not. Three pedagogical implications follow.

- **Supervision is a skill, and skills require practice.** A course that restricts AI to narrow, guardrailed uses, or that only shows students cleaned-up examples of AI at its best, does not prepare them to catch AI at its worst. Meaningful supervisory practice requires that students encounter AI in conditions where its failures are consequential — where they have to evaluate, challenge, and sometimes reject AI-generated work, and where the cost of not doing so is visible. The "AI as producer of artifacts to critique" and "AI as critic" patterns in §V.A speak directly to this: they cultivate evaluative judgment rather than operational fluency, and they do so at stakes the student has reason to care about.
- **Accountability belongs with the person, not the system.** Students should graduate understanding that using AI assistance does not transfer moral or professional responsibility to the tool. This is especially true in medicine, where all liability continues to rest with the clinician, regardless of AI tool use. This runs against the grain of how AI is often marketed — as a colleague, as an assistant, as a partner — and has to be taught explicitly. Disclosure norms, reflective practice built into assignments,

and modeling by faculty are three mechanisms; the AI posture statements in §V.A are the syllabus-level vehicle. The objective is not distrust of AI but calibrated trust: use it for what it does well, and keep the weight of judgment where it belongs.
- **The "illusion of competence" is the characteristic failure mode.** Students who produce good AI-assisted outputs often believe they have understood the underlying material, and are then surprised when they cannot reproduce or defend that work under conditions where AI is not available. This is both a learning problem (addressed by the meta-learning practices in §V.A) and an assessment problem (addressed by the proctored-assessment framework in §V.B). By the time the graduate reaches the workforce, it becomes a supervisory problem: a new hire who cannot tell when AI output is wrong is exactly the kind of junior contributor organizations will stop wanting to hire.

The point is not that AI use is dangerous. It is that AI use without developed supervisory judgment is dangerous in a way that specifically affects graduates, and that the university is the last structured venue with the time and the setting to build that judgment before the stakes become professional. This is what the employer expectation of "entering work ready to supervise, not merely use" implies on the educational side.

3. Guardrails for Engagement with Industry and Vendors

U-M's engagement with AI vendors and industry partners is already extensive: enterprise licensing arrangements, co-developed curricula, tool deployments inside Canvas, recruiting pipelines, advisory relationships with individual programs, and pilot arrangements with major providers. This engagement produces real benefits — tool access students could not otherwise afford, curricular relevance, placement opportunities — and the working group does not recommend reducing it. The question is what conditions engagement should satisfy to be defensible.

Six conditions are proposed, organized by the tension each addresses rather than as a generic list of principles. Each is intended to be auditable: an observer should be able to look at a specific partnership and say whether the condition is or is not being met.

- **Durational asymmetry.** Vendor product cycles, pricing structures, and feature sets change on timescales shorter than a degree program. Curricular commitments calibrated to a specific vendor tool — course assignments that presume a particular model's capabilities, assessment designs that rely on a specific interface — will outlive the product. The university should prefer to teach principles and postures that transfer across tools, and should treat vendor-specific training as an elective supplement rather than a core curricular spine. Where vendor-specific capability is unavoidable (certain professional programs, specific employer expectations), it should be scoped and named as such.
- **Vendor incentives are not institutional incentives.** AI tools offered free or at a steep discount to educational institutions are rarely philanthropy; they are customer acquisition. This is not a reason to refuse the offer, but it is a reason to account for the offer honestly — what the vendor gains (future customers, training data access, institutional endorsement) is part of the transaction, and the institution should negotiate with that in view. Faculty-survey responses describing AI adoption as vendor-driven and insufficiently consultative (Appendix F) are, among other things, a perception that these trade-offs have not been visible. Transparency about the terms of the deal is a legitimacy condition.
- **Pedagogical authority stays with faculty.** Partnerships must not drift into outsourcing educational decisions — what counts as meaningful learning, what constitutes mastery, how assessment is designed — to vendor platforms whose design optimizes for engagement and adoption rather than for learning. Platforms can be infrastructure; they should not become pedagogues. This is the same principle that animates the assessment framework in §V.B: integrity of what the university certifies depends on not outsourcing the certification.

- **Data sovereignty.** Student coursework, faculty intellectual output, patient information, and institutional data flowing through vendor-supplied AI represent a privacy and intellectual-property asymmetry that the institution is in a position to negotiate (see §V.D.2). Consent for training-data use, retention limits, and deletion rights should be contractual terms, not policy-page asides, and should be visible to students and faculty whose work is involved.
- **Equity and access.** Vendor relationships must not, in practice, create two-tier AI access across students or programs (see §V.D.1). Free tiers, student-discount tiers, and paid tiers produce real disparities in the capability students can bring to coursework; curricular expectations should be calibrated to the baseline the institution actually provides, not to what the best-resourced students can buy. Where premium capability is pedagogically warranted, the institution should provide it.
- **Reciprocity, not dependency.** Engagement with employers and vendors should be understood as a two-way channel in which the university also has something to say — about what counts as a durable education, about when a credential means what it says, about what it would take for a partnership to serve the student rather than only the partner. The institution's standing to make those claims is a function of how well the preceding five conditions are met; institutions that have not met them cannot credibly assert pedagogical authority against vendors that have identified real gaps in their practice.

These are not a checklist for vendor procurement. They are the terms under which partnership preserves rather than erodes the institution's integrity. The faculty-survey concerns about institutional AI messaging — that it has felt one-sided and vendor-adjacent — are best answered not by retreating from partnership but by being visibly transparent about the conditions each partnership satisfies.

4. Summary and Pointer to Recommendations

The three preceding subsections converge on a single institutional claim: U-M's comparative advantage is not speed of curricular response to industry; it is the cultivation, over the college years, of intellectual capacities that do not age on product cycles. Those capacities — framing, judgment, communication, supervision, and the ability to work with and through other people — are precisely what organizations increasingly need from graduates and increasingly lack the internal time to build. A curriculum that leans into that comparative advantage, rather than imitating industry training, is better-positioned on both educational and labor-market grounds.

The Recommendations section (§VI) proposes corresponding institutional investments: a preference for durable over vendor-specific framing in AI-related coursework; supervisory-practice design as an element of graduate-competency expectations (§V.C); procurement and partnership standards that make the six guardrails in §3 auditable rather than aspirational; and advisory structures that bring employer signal into the institution without letting it dictate the curriculum.

VI. Recommendations

The working group recommends that the University of Michigan approach generative AI through a spirit of deliberate experimentation, piloting diverse strategies, learning from what works, and scaling investment thoughtfully as evidence emerges. This iterative approach reflects U-M's commitment to rigorous inquiry applied to its own institutional practices.

Below, the working group offers its recommendations as follows. §1 covers the general **opportunity domains**. §2 frames the recommendations through **who** should have a role in realizing the recommendations as well as **what** they are, **where** they should be implemented, and **when**. §3 proposes the specific process for integrating **AI competencies in education**. §4 Summarizes the recommendations in a timeline.

1. Opportunity Domains

Instructional Opportunities

U-M is committed to ensuring that every graduate develops baseline AI competency. This competency is not defined by fluency with any particular tool, but grounded in conceptual understanding and ethical judgment. Central to this effort is meaningful support for faculty and students. U-M should **develop instructional design resources and AI literacy programming** to help instructors decide whether, when, and how to integrate AI tools confidently and critically. Student experiences will be carefully structured to balance AI-assisted work with hands-on, reflective learning that builds durable, AI-independent skills, and that will foster the development of judgment, which supersedes conversations about AI tools. In addition, students should develop an understanding of AI as a sociotechnical system shaped by human choices, biases, and values; using AI productively as a collaborative learning partner informed by genuine domain expertise; and maintaining full accountability for one's work, including appropriate disclosure and attribution of AI assistance. Importantly, schools, colleges, and units will retain the autonomy to tailor these competencies to their disciplinary contexts, supported by shared university resources and expertise. We propose a process for integrating AI competencies in the next section.

The working group recognizes that **responsible AI integration does not mean universal AI adoption**. There have been calls for AI-integrated, AI-limited, and AI-free pedagogical environments. The working group suggests that they be investigated further, starting with low-cost experiments (e.g., social spaces that signage designates as “AI free”). These experiments could be piloted on a small scale as voluntary, socially-normed spaces rather than technologically enforced restrictions and would offer faculty and students room to explore the value of unmediated learning. The committee unanimously recommends that students should be clearly informed about the role of AI in any course, including relevant expectations and tools, communicated through the syllabus and course materials.

Assessment Opportunities

The working group recommends that **U-M not embrace surveillance-based AI detection tools**, which have proven unreliable and counterproductive. Instead, institutional policy should center on behavior-based language, addressing the misrepresentation of authorship and the outsourcing of intellectual work. The working group recommends **revisiting Honor Code practices around proctored exams**.

The university should **invest in a broader assessment infrastructure**. This includes proctored testing environments, oral exams, in-class demonstrations, and other interactive formats, with the scale and form of investment to be determined by piloting at a manageable level before committing to major infrastructure. Schools and colleges will retain authority over the assessment approaches best suited to their disciplines. The working group recommends that U-M include undergraduate, graduate students and SSD in this effort. U-M should consider commissioning student cohorts to help design and evaluate assessment strategies, an approach that deepens student engagement while potentially cultivating interest in graduate and research pathways.

Infrastructure Opportunities

The working group recommends that U-M should stand up two advisory boards: an **internal Academic-AI Advisory Board** and an **External-AI Advisory Board**. These advisory boards will ensure that governance of AI in learning reflects broad perspectives and is not concentrated within any single administrative unit. Decisions, wherever possible, will be made at the unit-level, empowering the communities closest to the work. The internal board would bring together faculty, students, staff, and academic leaders from across schools, colleges, and administrative units to monitor the evolving AI landscape, assess which initiatives are gaining traction and which are falling short, and make informed recommendations about where to invest further and where to pull back. Critically, this body would serve as a counterbalance to the perception that technology decisions are made solely within ITS, ensuring that academic priorities, pedagogical values, and disciplinary diversity remain at the center of AI governance. The internal board would also serve a coordinating function, working alongside AI-related committees that individual schools and colleges may establish, fostering coherence without imposing uniformity. The external board would provide independent perspective from outside the institution, drawing on expertise from peer universities, industry, policy, civil society, and communities affected by AI, helping U-M remain accountable to broader public interests and stay attuned to developments beyond its own walls.

U-M should continue to invest in the **technical infrastructure** necessary for experimentation (e.g., sandbox environments and API access). This will allow students and faculty to explore, build, and critically evaluate AI tools firsthand and will begin to address equity issues stemming from lack of access to various types of AI technologies.

2. The Roles, Content, Location, and Timing of the Recommendations

Who needs to evaluate and/or implement the integration of AI into teaching and learning?

According to the MIDAS-commissioned benchmarking study, while institutions are responding to AI in education, their policies, tools, coursework, and requirements vary widely. Universal undergraduate AI literacy requirements are rare, competency frameworks are sharply split, and formal university-level assessment is almost nonexistent. U-M can lead AI education by pairing institution-wide, ethics-centered AI literacy outcomes with equitable access to vetted tools and transparent assessment. AI education at U-M is comprehensive in scope, involving students, faculty, and staff.

- **Students:** To be leaders in their respective fields, U-M students need to graduate with the ability to understand and use AI competently, having gained skills and familiarity with AI tools through university exposure. Students need to learn to use AI when suitable, choosing appropriate roles for AI (e.g., explainer, coach, critic, simulator) and judging each use by its likely benefits, risks, and the design choices that make it educationally effective.

Recommendation: As AI becomes embedded across academic, creative, and professional settings, U-M must equip students to understand how AI functions, how to use AI responsibly, and how to evaluate AI results/outcomes.

- **Faculty:** According to the faculty survey, 85% of faculty have already used generative AI to some degree. As AI use has increased and developed across disciplines, U-M faculty have demonstrated that they are ready to provide students with the skills they may need upon graduating by incorporating AI use into classes.

  Recommendation: Curriculum committees within each department, program, or college should determine how, where and if AI fits appropriately into their courses. Faculty teaching these courses should have support to experiment, make pedagogical changes, and run pilot studies to explore AI use in their classes. If faculty decide not to use AI in their classes, they should also be supported within the unit, and the AI use incorporated elsewhere in the degree coursework as appropriate.

- **Staff:** Any institution expecting to be at the forefront of leading a bold initiative such as developing competency in AI requires examining all aspects of the university, including assessment of staff needs and training for the use of AI.

  Recommendation: Departments/Programs should evaluate what AI use and training may be needed for staff.

What defines AI competency?

Faced with the complex charge to define AI competency, the working group identified two key, distinct aspects of AI education:

1. **AI for education**: In the preceding sections of this report, the working group summarized a learning-centered framework for evaluating and designing AI use in courses by weighing potential benefits, potential harms, and intentional design choices in discipline-appropriate adoption.

   Recommendations: Instructors should outline concise expectations of AI use as deemed appropriate for the course. Since AI use will vary greatly in each discipline, deployment of AI tools in coursework should be evaluated and implemented by program curriculum committees. However, some general recommendations could be made for instructors to clarify/outline AI use in their course. Examples for how to accomplish this goal could be by 1) embedding the benefits/harms/design-choices framework in the course, either through a student activity applying it to course-relevant AI uses, or 2) via an “AI posture” syllabus statement explaining the instructor’s choices and rationale.

2. **Education for AI:** Students should be introduced to AI as something people build, running through many steps from gathering data to putting AI to use as a tool, where human decisions shape what it does and who it helps/harms. It’s important for students to learn how AI outputs are best understood as likely guesses rather than verified facts since the same prompt can produce different answers, and responses can sound confident while still being wrong.

Recommendation: Students should develop AI literacy by actively guiding, evaluating, and potentially pushing back against AI systems in their roles as citizens, learners, and professionals. Reinforcing this literacy is the responsibility of all instructors who incorporate AI into their courses.

Where should AI be incorporated into the curriculum?

The working group emphasizes that there is no "one size fits all" approach to how AI should be incorporated into curriculum across all disciplines. Each college, department, or program knows the needs of its students best, and faculty within these units are versed in the use of specific AI tools that relate to their fields. Many units already have instructors using AI tools in course curricula and could serve as mentors for faculty newly incorporating AI use in classes.

Faculty will need additional support to enact these coursework changes. Undergraduate Learning Experience Designers (LXDs) who are AI-fluent could help instructors implement advanced course-design alterations without each instructor starting anew.

Recommendation: Curriculum committees within units should assess what AI skills are needed for their students and determine which courses within the degree program could be suited for (or are already) implementing 1) the deployment of AI tools and 2) reinforcing AI literacy. The AI in Education Working Group recommends no credit hour changes occur from incorporating AI skill development in coursework. To support course changes, additional trained staff or undergraduate students will be needed.

When should changes begin?

- U-M's expectation is that graduates become active, responsible AI users who can understand and critically evaluate AI systems, therefore changes should begin immediately to best serve our graduates. Entry-level readiness now means not just performing tasks, but supervising AI effectively. U-M should prioritize long-lasting skills that include problem framing, judgment, communication, supervision, and collaboration, over keeping pace with the latest industry tools.
- Building university-wide support for thoughtful, ethical AI use in learning to develop baseline competencies, provide resources for faculty and students, and enable experimentation in AI use should be a phased process. Due to the rapidly changing nature of AI in both industry and academia, only the first two years are outlined here.

3. Proposed Process for the Integration of AI Competencies

The AI in Education Working Group developed a phased plan to integrate discipline-specific competencies across the schools and colleges of U-M.

Phase 1 — Provost-Level Mandate and Vision

By an overwhelming majority U-M faculty voted in November 2025 to support Motion 4, a "Resolution Demanding a Comprehensive Policy Addressing the Role of Generative Artificial Intelligence at the University of Michigan." The Provost and the Vice-President for Information Technology constituted a working group, AI in Education Working Group. The group was tasked with exploring the thoughtful, ethical, and practical integration of artificial intelligence into teaching and learning at the University of Michigan. This included the relationship of AI to models of liberal arts inquiry and instruction, as well as circumstances in which AI may not be appropriate for use based on desired pedagogy, assessment criteria, and learning outcomes.

Specifically, members of the working group were asked to:
- Review the prior U-M report on generative AI sections on teaching and learning to assess progress and relevance;
- Evaluate the current status of AI use in teaching and learning, and identify future needs and opportunities.
- Identify leading practices and models in AI-based teaching and learning from within U-M and peer institutions.
- Focus exclusively on the teaching and learning applications and impacts of AI, distinct from research or university business processes.
- Provide guidelines for evolving pedagogy and assessment in the age of AI.

Phase 2 — Cross-College Faculty Task Force

The AI in Education Working group was created. It was a multidisciplinary task force composed of key faculty, researchers, and administrative leaders from LSA, Michigan Medicine, the College of Engineering, Ross Business School, the Law School, College of Pharmacy, School of Information as well as from the UM-Dearborn and Flint campuses. This report is the outcome of that task force. It includes the production of a shared AI Competency Framework, a set of 3–5 principles that all colleges must satisfy, without specifying how each must do so. This is the "what," not the "how."

Phase 3 — College-Level Curriculum Design

We propose that this document forms the beginning of a shared infrastructure. From here, each college's Curriculum Committee will design its own discipline-specific competencies and delivery model. The key design constraint: no added credit hours, and competency must be embedded into existing degree requirements.

Examples from Schools and Colleges:
- *Engineering*: The new AI minor in Computer Science and Engineering can serve as the backbone; existing courses can be designated as satisfying the competency
- *School of Information*: The human-centered AI minor provides an ethics-and-risk track that other colleges could cross-list
- *Ross*: Ross graduate programs already permit AI use with required citation; an undergraduate competency module can build on existing business analytics courses
- *College of Pharmacy*: Current, required undergraduate coursework in Medicinal Chemistry for the BSPS degree introduces students to AI- and data-driven computational workflows to explore applications in drug discovery and development.

Key Course Design Principles for U-M
- *Lead with faculty*. The university should support faculty innovators in AI design, development, and pedagogy. Faculty should be involved in meaningful ways in the development of all AI policies, pilots and task forces, including consultation with faculty governance.
- *Build on what exists*. U-M already has the AI infrastructure (U-M GPT, Go Blue), the minors and coursework (Engineering CSE AI minor, School of Information human-centered AI minor, PCAS), and the governance precedents.
- *Distinguish technical from ethical competency*. The framework should require both dimensions, with colleges determining the appropriate ratio for their discipline.

Phase 4 — Pilot Cohort (Fall 2026 – Spring 2027)

Rather than launching fully before standards are proven, U-M should run a voluntary pilot across 2–3 colleges for the first year. Engineering and the School of Information are the natural pilots given their existing AI minors. The pilot generates real assessment data, student and employer feedback, and surfaces implementation problems before the requirement is mandatory.

Phase 5 — University-Wide Launch (Fall 2027)

The requirement applies to all incoming undergraduates. Competency is demonstrated through discipline-specific projects embedded in existing coursework — not a standalone exam or a new required course. Colleges may choose to designate existing courses as fulfilling the requirement, create new competency modules within existing courses, or develop new shared electives (especially for smaller colleges without existing AI offerings).

Phase 6 — Annual Refresh Cycle

Each college establishes a standing Advisory Board focused on AI competency needs of their graduates for careers and future education. Boards meet annually, report to the dean and curriculum committee, and trigger a lightweight review of whether the competency standards need updating. Outcomes data (employer surveys, graduate placement, student self-assessment) feeds back into Phase 3 continuously.

4. Summary of Recommendations Timeline

2026 – 2027

- Create an internal Academic-AI Advisory Board
- Establish faculty and student support by providing instructional design/AI literacy resources to upskill instructors.
- Encourage faculty to create structured student experiences that balance AI-supported work with hands-on, reflective, AI-independent learning.
- Enable schools, colleges, and units to tailor competencies and responsible use to field-specific goals, supported by shared resources and experts.
- Provide sandbox environments and API access so students and faculty can experiment.
- Support diverse pedagogy by acknowledging and supporting AI-integrated, AI-limited, and AI-free environments, including certified AI-free workspaces or digital-free zones.
- Support pilot projects that examine student competency, such as proctored testing centers (using existing facilities), providing tutoring centers, or providing resources for oral exams.

2027 – 2028:

- Create an External-AI Advisory Board.
- Provide additional sandbox environments and API access so students and faculty can experiment.
- Review student competency assessments to determine if students' work reflects their own abilities.
- Establish guardrails so partnerships and tools support its academic mission and public responsibilities as industry influences grow.

- Assess infrastructure needs and invest in proctored testing centers, expanding proctored assessment options (including low/AI-free spaces).
- Expand resources for interactive assessments such as oral exams to confirm student competencies.

VII. Implications and Future Considerations

In this section, we first discuss risks of inadequate action or overly prescribed strategies. Then we present a number of further considerations and recommendations to support a deeper, more thorough and more forward-looking institutional strategy.

Risks of inaction or overly prescriptive approaches. The working group believes that a balanced approach for AI's use in education and AI education is essential.

As all higher education institutions plan and implement AI strategies, inaction or inadequate action will jeopardize our university's leading role in higher education. Inaction could include the lack of rigorous thinking, the lack of evidence-based planning, and the lack of practical support. Instructors and students are all considering or experimenting with how to use AI to support teaching and learning. Inaction at the institutional level means that instructors and students will be largely on their own, which will result in inconsistency in both conceptual and practical approaches. Instructors will lack guidance in thinking about how to use AI to support teaching and learning, both when they decide to adopt AI and not to adopt AI. In addition, with institutional inaction, much of students' AI use will develop (and is already developing) informally through peer norms. Institutional policy and support, including formal instruction and faculty guidance, can play an essential role in shaping how students interact with AI and think about the role of AI in their study, career and lives at the critical formative stage. Therefore, it is all the more important for the university to engage explicitly with students about appropriate and effective use of AI, instead of inaction. This is also important because universities should be involved, more broadly, in addressing emerging regulatory, ethical and liability risks of AI and in policy development. Equally importantly, AI literacy and the ability to appropriately use AI in many types of work will be demanded by employers across all lines of work. The lack of institutional strategies for AI literacy and even proficiency for all of our students means that our students will be negatively impacted on the job market, have less flexibility in professional pursuits, and miss the opportunities to lead how AI shapes their professions.

On the other hand, over-prescription will be almost equally counter-productive for several reasons. First, AI's role in teaching and learning, while sharing commonalities across fields of study, will likely look differently for each field. For example, how computer science students learn programming and how history students build their perspectives will need to leverage AI in very different ways. Computer science students will learn how AI can be an effective collaborator in coding, while history students will still need to read all the texts on their own. It is not feasible to have an institution-wide AI strategy for education that is nuanced enough for all fields. Rather, instructors and students in each field will need to leverage institutional guidelines and infrastructure to fit with, and transform, the norms in their own field. Second, AI should not replace human individuality. Students come with different learning goals, styles and preparations. Instructors infuse their own creativity and perspectives in their teaching. Each encounter between an instructor and a student is unique. Such individualized experiences educational attainments should be preserved, not weakened, by AI. AI's role in flattening perspectives and human expressions is already alarming. We should use AI to support individualized teaching and learning, not to homogenize it.

Topics for further exploration. The working group's recommendations are based on our current understanding of AI and our current thoughts on the purposes of undergraduate and graduate educations, and are focused on immediate actions. This is the only starting point that is possible for us. Hence, our analyses and recommendations are based on limited understanding and evidence of both AI's capabilities and impacts, and on how educators and students reevaluate the purposes of education. Here we list a few topics, among many, that require further exploration.

AI-native students. The current undergraduate and graduate students are exploring AI tools, so are the faculty. However, very soon (if this hasn't already happened), students will enter college already having a lot of exposure to AI because AI is already being implemented in K-12 education, and because many students have easy access to AI in daily life. They will be the AI-native students who already make AI an indispensable part of their lives. Instructors will face new challenges in guiding their AI use and determining what is allowed or not allowed.

Expectations from employers. Universities are partly preparing students for jobs. AI literacy and proficiency will inevitably become more essential in our graduates' employability. However, employers' expectations of AI-augmented graduates are shifting rapidly and inconsistently. How can the university stay calibrated to a moving target?

Beyond faculty and students. Most of the current AI-in-education discussions focus on faculty and students, but many other stakeholders and contributors to education should be included. We cannot meaningfully adjust educational methods without ensuring that administrative systems, staff, academic advisors, parents, and stakeholders in both the private and public sectors, whose main role is not necessarily in the classroom, are an integral part of the conversation and preparation. This also includes upskilling staff when we plan for upskilling faculty, and supporting autonomy for staff as well as for faculty.

Equal access. AI access, literacy, and confidence are not evenly distributed across socioeconomic backgrounds, disciplines, or geographies. For example, first-year undergraduate students will vary greatly in their exposure to AI and their skills to use AI appropriately to support their learning goals. Policies and practical support that assume uniform access will inadvertently widen existing gaps. Instead, we need to actively support equal access.

"AI in context." Norms of fields of study will change quickly as AI enters each field in different ways and warrants different considerations on how AI should shape each field. As such, how we should use AI to teach students in each field, how we should define AI competency for each field, and how we should prepare students so that they have the flexibility to enter other fields, will need to be a continuous process.

An evidence-based strategy for AI in higher education. Research on the impact of AI on educational practices and student outcomes is only starting and results are sparse. But such research will be critical as universities refine their AI strategies. In the coming years, the university should not only support such research, but also use such research to guide the evolution of institutional policy, support mechanisms, and teaching and learning practices. For example, if we worry about the negative impact of AI on student cognitive abilities, we need to understand what impact is happening and why, and how to use AI to enhance cognitive abilities. Same with other concerns such as the loss of individuality in the learning process.

Relatedly, we also need to establish an environment that encourages experimentation with AI (a "sandbox" environment). Given the need for instructors to experiment with effective ways to incorporate AI in teaching and learning in order to achieve the desired outcomes, and the emerging regulatory and liability risks associated with AI, the "sandbox" can both help limit potential regulatory and liability risks and allow trial and error.

Anticipated advancement of AI. A challenge that is unique to AI, as compared to other learning technologies, is the fact that AI is evolving at a pace that is unseen with other technologies. Its pace of advancement is much faster than universities traditionally update curricula and policy, let alone any significant reconsideration about what counts as essential intellectual growth and practical skill development for the future. The bigger challenge, beyond what we recommend for the immediate future at the practical level, is to focus on the "North Star": the purposes of education, and use that to assess how upcoming AI tools align with such purposes, and to develop ways to help our community continuously adapt.

Beyond the immediate support, we should also make room (time, facilitation, and coordination) for our community to continuously explore the purposes of higher education and determine the essential qualities that students need to develop. These will be independent of specific capabilities of new AI systems. On the other hand, specific skills and conceptual capabilities may well evolve alongside AI evolvement, and we need to provide long-term support for our community to gain a basic understanding of new AI systems (what they are capable of doing, how they interact with humans, how much we can trust them), so that they can determine what roles these new AI systems can play in higher education and shape the desired outcomes.

Appendices

Appendix A: Working Group Membership

Mika LaVaque-Manty, Chair
Arthur F. Thurnau Professor, Associate Professor of Political Science and
Associate Professor of Philosophy, College of Literature, Science, and the Arts

Laura Clifford
Lecturer III in Medicinal Chemistry, College of Pharmacy

Ivo Dinov
Henry P. Tappan Collegiate Professor of Nursing, Professor of Nursing
Director Academic Program, School of Nursing
Professor of Computational Medicine and Bioinformatics, Medical School

Vikramaditya Khanna
William W. Cook Professor of Law and Professor of Law, Law School

Jing Liu
Executive Director, Michigan Institute for Data and AI in Society

Emily Mower Provost
Professor of Electrical Engineering and Computer Science
Associate Chair, Division of Computer Science and Engineering,
Department of Electrical Engineering and Computer Science,
College of Engineering and Professor of Psychiatry, Medical School

Venkatram Ramaswamy
Professor of Marketing, Computer and Information Systems,
Stephen M Ross School of Business

Paul Resnick
Michael D Cohen Collegiate Professor of Information
Professor of Information, School of Information

Nadja Rottner
Associate Professor of Art History, Dept of Language, Culture, and the Arts,
College of Arts, Sciences and Letters, University of Michigan-Dearborn

Yael Sela
Collegiate Lecturer, Lecturer IV in Psychology, Dept of Psychology
Chair, Department of Psychology, College of Arts, Sciences and Education, University of Michigan-Flint

LaKisha Michelle Simmons

Arthur F. Thurnau Professor, Associate Professor of Women's and Gender Studies Associate Professor of History, College of Literature, Science, and the Arts

Max Spadafore
Director of Educational Informatics for the Office of Medical Student Education
Clinical Assistant Professor of Emergency Medicine, Medical School

Yuekai Sun
Associate Professor of Statistics and Director Academic Program, Statistics
College of Literature, Science, and the Arts

Staff Support
Zarinah Aquil, Project Senior Manager, Office of the Provost
Mallory Jane "MJ" Kwiatkowski, IT Project Manager Senior, Information and Technology Services

Appendix B: Charge and Scope Document

Revised 1/22/2026

Co-Executive Sponsors

- **Laurie K. McCauley**, Provost and Executive Vice President for Academic Affairs
- **Ravi Pendse**, Vice President for Information Technology and Chief Information Officer

Statement of Purpose

The AI in Education Working Group is established to explore the thoughtful, ethical, and effective integration of artificial intelligence into teaching and learning at the University of Michigan. This exploration will include the relationship of AI to models of liberal arts inquiry and instruction, as well as circumstances in which AI may not be appropriate for use based on desired learning outcomes. The Working Group will consider balancing AI usage with maintaining core knowledge in the fundamentals of the discipline of study and promoting critical thinking skills. The group will build on prior working group reports by reviewing their findings, assessing current initiatives, and identifying future opportunities to strengthen the educational impact of AI across the university.

Objectives

The primary objectives of the working group are:

- **Review** prior U-M report on generative AI to assess progress and relevance.
- **Evaluate** the current status of AI use in teaching and learning, and identify future needs and opportunities.
- **Identify** leading practices and models in AI-enhanced education from within U-M and peer institutions.
- **Focus** exclusively on the educational applications and impacts of AI, distinct from research or business uses.

Scope

The working group's scope focuses exclusively on the educational applications and impacts of artificial intelligence across **U-M's Ann Arbor (including Michigan Medicine), Dearborn, and Flint campuses**. This includes applications that enhance pedagogy, curriculum design, student learning, assessment, and instructional support. The group's charge does not extend to AI use in research activities or university business processes.

Working Group Key Deliverables

- Comment on June, 2023 Generative Artificial Intelligence Committee Report
- Understand AI practices and policies already underway across the university, focusing on educational uses
- Conduct benchmarking among peer institutions to identify exemplars
- Identify best practices for use of AI in the educational context
- Articulate (at a macro-level) the AI competency UM graduates should possess, informed in part by stakeholder consultation with prospective employers of U-M graduates.
- Recommend a framework for coordinating university-level and unit-level approaches to AI

Timeline

- Working Group to be charged in December, 2025
- Report due by April 30, 2026

Initial Materials Provided

1. Generative Artificial Intelligence Committee Report - June 2023
2. Artificial Intelligence Research Committee Current State Report (Office of the Vice President for Research) - December 2023 *(Confidential Internal Report, Not Intended for Public Distribution)*
3. Artificial Intelligence Research Committee Recommendations Report (OVPR) - April 2024 *(Confidential Internal Report, Not Intended for Public Distribution)*
4. High-Performance Computing Committee Recommendations Report - August 2024 *(Confidential Internal Report, Not Intended for Public Distribution)*
5. AI Institute Working Group Report - December 2024 *(Confidential Internal Report, Not Intended for Public Distribution)*
6. Generative AI and Education Digital Pedagogies, Teaching Innovation and Learning Design (2024)

Appendix C: Meeting Summaries

Meeting #1

- **Date**: December 12, 2025
- **Focus**: Establishing the working group's charge and scope, specifically focusing on teaching, learning, and pedagogy related to Generative AI at U-M.
- **Key discussion themes**:
  - Identification of gaps in current AI guidance for faculty and students.
  - Importance of defining baseline AI competencies for all U-M graduates.
  - Addressing faculty resistance and the need for discipline-specific differentiation.
  - Consideration of the environmental, social, and ethical implications of AI.

Meeting #2

- **Date**: January 16, 2026
- **Focus**: Refining the core domains for the final report and discussing the definition of AI "competency" versus "proficiency".
- **Key discussion themes**:
  - Emphasis on "understanding under the hood" as a core competency; how AI works, not just how to prompt it.
  - The need for U-M students to be able to appraise and validate AI models for specific use cases.
  - Inclusion of ethics, equity, and access as central components of the report.
  - Supporting students who choose to refuse AI use based on personal ethics.

Meeting #5

- **Date**: February 23, 2026
- **Focus**: Primary discussion on assessment strategies and reviewing the scope of the upcoming faculty survey.
- **Key discussion themes**:
  - The front-line challenges and low morale of GSIs managing AI-related academic integrity issues.
  - Specific concerns within the Humanities, including the "vulnerability and opportunity lens" of AI.
  - Planning for a benchmarking study of approximately 100 peer institutions.
  - Teaching students how to explicitly use AI as a tutor for learning.

Meeting #6 (Listening Session)

- **Date**: March 13, 2026
- **Focus**: Presentations from faculty on academic integrity policies and the practical use of RAG-based AI "tutors" and virtual patients.
- **Key discussion themes**:
  - Focusing academic integrity policies on student behavior and authorship rather than specific tools.
  - The effectiveness of personalized AI tools like "Wolverine Tutor" in improving learning outcomes.

  - The tension between data privacy and the pedagogical need for faculty to see student-AI interactions.
  - Framing AI as a "playground" for students to understand future workforce impacts.

Meeting #8

- **Date**: March 30, 2026
- **Focus**: Reviewing preliminary faculty survey results and drafting short-term recommendations for the 2026-27 academic year.
- **Key discussion themes**:
  - Reviewed initial faculty survey data.
  - The perceived "pervasive and undetectable" nature of student AI use versus actual student usage data.
  - Recommendation for "thousand flowers" pilot programs to experiment with new assessment models.
  - The necessity of hands-on faculty support and undergraduate consultants to bridge the AI expertise gap.

Meeting #9 (Listening Session)

- **Date**: April 17, 2026
- **Focus**: Debriefing benchmarking data and finalizing the framing and core themes for the final report due on April 30.
- **Key discussion themes**:
  - Benchmarking results showing U-M is "medium and average" compared to R1 peers, with a notable gap in competency frameworks.
  - A commitment to framing the report beyond the "pro vs. against" binary to focus on shared threats and commonalities.
  - A listening session regarding "vibe coding" and using AI to teach ethical hacking.
  - Proposal to integrate AI literacy into First-Year Writing Requirement (FYWR) courses.

Meeting #10

- **Date:** April 27, 2026
- **Focus:** Finalizing the structural decisions and thematic content for the final report due April 30, with an emphasis on implementation pathways and executive framing.
- **Key discussion themes:**
  - Report Structure and Tone: Deciding to integrate "legitimacy" into the framing rather than a standalone section and intentionally preserving variation in tone to reflect human authorship.
  - Equity and Access: Addressing systemic risks including racial bias in LLMs, environmental impacts, and "lock-in" pricing strategies from vendors that could disadvantage students.
  - Assessment Strategy: Moving away from unsustainable detection and "policing" tools toward rethinking assessment models that prioritize the learning process.
  - Workforce Readiness: Shifting the focus from technical execution to "durable capacities" like judgment and accountability as entry-level roles transition into supervising AI agents.
  - Executive Framing: Highlighting the inherent tension between AI's goal of "reducing friction" and the "productive friction" required for deep, transformative learning.

Appendix D: Materials Reviewed

University of Michigan Internal Reports & Documents

These materials were identified as the foundation for the AI in Education Working Group's research and scope:

- **Generative Artificial Intelligence Committee Report (June 2023)**
- **OVPR AI Reports** (Office of the Vice President for Research):
  - Artificial Intelligence Research Committee Current State Report (December 2023) - *Confidential internal use only*
  - Artificial Intelligence Research Committee Recommendations Report (April 2024) - *Confidential internal use only*
- **High-Performance Computing Committee Recommendations Report (August 2024)** - *Confidential internal use only*
- **AI Institute Working Group Report (December 2024)** - *Confidential internal use only*

Academic Frameworks & Pedagogical Models

The curated list and meeting notes reference several established educational theories adapted for the AI era:

- **SAMR Model**: Used for scaffolding competence (Substitution, Augmentation, Modification, Redefinition).
- **TPACK (Technological Pedagogical Content Knowledge)**: A framework for teacher knowledge required to integrate technology.
- **ABC Learning Design Framework**: Based on Laurillard's Conversational Framework, used to convert face-to-face learning to online/AI delivery.
- **Bloom's Taxonomy**: Specifically referenced in the context of shifting from version 2.0 to 4.0 for AI-age objectives.
- **Generative Learning Theory**: An instructional design framework predating the digital age but applied to GAI.
- **Community of Inquiry (CoI) Framework**: Focused on social, cognitive, and teaching presence.

Curated AI Tools & Platforms

The following tools are discussed as exemplars or active resources:

- **U-M Specific Tools**:
  - **U-M GPT**: A private, accessible alternative to OpenAI's ChatGPT.
  - **U-M Maizey**: A no-code platform for building custom RAG-based (Retrieval-Augmented Generation) chat programs.
  - **U-M GPT Toolkit**: Advanced tools for staff and student use.
  - **Wolverine Tutor**: A specific Maizey project used to help students practice towards mastery.
- **External AI Platforms**:
  - **Perplexity, Elicit, Semantic Search, and Scite**: Tools for targeted research and investigation.
  - **Khan Academy (Khanmigo)**: Referenced for experiments in prompt engineering and lesson planning.

    - **Oak National Academy**: Used as an example of implementing Retrieval-Augmented Generation (RAG) to link AI output to curriculum standards.
    - **Circle Chat**: Designed for group interaction with AI avatars.

External Literature & Key References

Notable books and articles cited as influential to the working group's framing:

- European Commission: Directorate-General for Education, Youth, Sport and Culture. *Guidelines on the Ethical Use of Artificial Intelligence and Data in Teaching and Learning for Educators*. Publications Office of the European Union, 2026, https://data.europa.eu/doi/10.2766/7967834.
- Koehler, M. J., and P. Mishra. "Introducing TPCK: An Adoption of Technology Pedagogical Content Knowledge Framework." *AERA 2008*, 2008.
- ---. "What Is Technological Pedagogical Content Knowledge?" *Contemporary Issues in Technology and Teacher Education*, vol. 9, no. 1, 2009, pp. 60–70.
- Mollick, Ethan. *Co-Intelligence: Living and Working with AI*. Portfolio, 2024.
- Pratschke, B. Mairéad. *Generative AI and Education: Digital Pedagogies, Teaching Innovation and Learning Design*. Springer Cham, 2024, https://doi.org/10.1007/978-3-031-67991-9.
- Suleyman, Mustafa, and Michael Bhaskar. *The Coming Wave: Technology, Power, and the Twenty-First Century's Greatest Dilemma*. Crown, 2023.
- UNESCO. *Guidance for Generative AI in Education and Research*. By Fengchun Miao and Wayne Holmes, UNESCO, 2023, https://doi.org/10.54675/EWZM9535.
- Warr, M., and P. Mishra. "TPACK: The TPACK Technology Integration Framework." *EdTechnica: The Open Encyclopedia of Educational Technology*, EdTech Books, 2022, https://doi.org/10.59668/371.9034.

Appendix E: Consultations and Informational Conversations

The working group held several listening sessions with faculty from across the university—including representatives from the Law School, Michigan Medicine, Sweetland Writing Center, and the U-M School of Information—to gather diverse perspectives on AI's role in the classroom.

Listening Sessions

Purpose of Engagement

The Working Group invited instructors from across the university to share their experiences with generative AI. These sessions allowed the group to document the wide variation in faculty sentiment and practical AI adoption.

Specifically, the group sought to:

- **Understand pedagogical experimentation**: Examine real-world use cases, such as "virtual patients" for medical training and AI clinics for law students, to identify what is working and what challenges persist.
- **Identify institutional gaps**: Surface specific needs for technical infrastructure, pedagogical guidance, and university-level policy.

- **Address concerns**: Hear directly from faculty who are skeptical or opposed to AI to ensure the final report acknowledges the risks to critical thinking and academic integrity.

How Input Informed the working group's Thinking

- **Validating the "Split Campus"**: Hearing from both critics and early adopters led the group to frame recommendations that support diverse approaches, ensuring that faculty who choose AI-free coursework feel as supported as those who integrate the technology.
- **Shifting Policy from Tools to Behavior**: Faculty input encouraged the group to recommend behavior-based academic integrity guidelines (focusing on authorship and disclosure) rather than trying to regulate specific, rapidly changing tools.
- **Prioritizing Process Over Product**: Demonstrations of "vibe coding" and "ethical hacking" shifted the group's focus toward grading the process of learning, such as the quality of student troubleshooting and prompting, rather than just the final output.
- **Technical Requirements for ITS**: Feedback from instructors on the limitations of current campus tools led to specific recommendations for future features, such as the ability for students to export their AI conversation histories to instructors for assessment and feedback.

Director of Sweetland Center for Writing

Purpose of Engagement

The working group met with the Director of Sweetland to discuss integrating **AI literacy into the First-Year Writing Requirement (FYWR)**. This was identified as a strategic way to reach a large portion of the student body through a common syllabus framework while maintaining independent instructor syllabi.

How Input Informed the Working Group's Thinking

The Director's input led the WG to view Sweetland's summer work on AI as a potential **framework for university-wide AI competencies**. This moved the group toward recommending a phased approach to student literacy, starting with basic AI understanding in foundational courses.

AAC&U Team (Matt Kaplan, Pete Bodary)

Purpose of Engagement

These individuals provided insights into national trends and local practical applications. Matt Kaplan (CRLT) shared information on **national toolkits for AI teaching and learning** through the American Association of Colleges and Universities (AAC&U). Pete Bodary presented on his use of "Wolverine Tutor" and "virtual patients" using U-M's Retrieval-Augmented Generation (RAG)-based Maizey system.

How Input Informed the Working Group's Thinking

- **Kaplan's** input helped the working group benchmark U-M against national standards, noting that while U-M has strong tools, it lacks a formal competency framework compared to peers.
- **Bodary's** presentation shifted the working group's focus toward **process-oriented learning**. The group discussed a recommendation to curate "thoughtful use cases" like his to demonstrate how AI can be a "Socratic tutor" rather than just a shortcut as a resource for faculty.

Provost's Faculty Advisory Committee (PFAC)

Purpose of Engagement

The working group utilized PFAC as a key channel to gather broad faculty sentiment and understand concerns from those skeptical of AI. The goal was to ensure the report represented the "wide variability in faculty familiarity, comfort, and willingness to engage with AI".

How Input Informed the Working Group's Thinking

Feedback from these representatives underscored the **"vocal minority" of faculty** who feel unheard and strongly opposed to AI. This led the working group to adopt a "neutral" framing for the final report—moving away from a "pro- or anti-AI" stance—and ensuring recommendations explicitly support **AI-free classrooms** alongside AI-integrated ones.

CRLT & LSA - Departmental Conversations / Workshops

Purpose of Engagement

These sessions aimed to understand discipline-specific barriers and the practical workload faculty face when redesigning assessments. The group specifically looked at how the humanities disciplines and the U-M College of Literature, Science and the Arts (LSA) were navigating the potential loss of "productive struggle" in writing and analysis. Conversations were held with the Center for Research on Learning and Teaching (CRLT) to understand current offerings to assist faculty, and future plans.

How Input Informed the Working Group's Thinking

These departmental insights highlighted that **"do what you want" is insufficient guidance** for faculty. Consequently, the working group prioritized recommendations for **scalable, hands-on support models**, such as undergraduate "Blue Corps" consultants, to help faculty upskill without adding to their existing research and teaching burdens.

Appendix F: Faculty Survey Results

Methodology

The survey was designed for "planned missingness," to reduce the time it would take any given single respondent to complete it. This was achieved through two techniques:

1. **Role-based logic:** Respondents' early answers to the roles they had (e.g., teaching, research) were used to give them only relevant questions.
2. **Modular randomization:** Only subsets of large question banks on any given topic were given to a single respondent. The combination of randomization and total numbers answered nevertheless guaranteed a sufficient sample of answers. Inferences to the entire set of respondents were then made by imputation, i.e., estimating how the subsets would likely have been answered based on what had been answered.

The survey was distributed on March 16-17 to 11,630 University of Michigan faculty on all campuses (Ann Arbor including Michigan Medicine, Dearborn, and Flint), including research and clinical faculty, librarians and archivists with faculty appointments, adjunct faculty, as well as faculty with emeritus status. It closed on March 27, although it remained open and late answers were accepted until April 5. Faculty positions and units were linked into the data from the distribution lists, but no direct identifiers were. These elements were disaggregated in the analysis in such a way that identification of an individual due, e.g., to their multiple appointments was not possible. Here and in future reporting of results, categorizations that might lead to the identification of individuals are omitted.

The survey was designed by Professor Josh Pasek, with graduate student research assistants Cydney Grannan and Amanda McKeever, and undergraduate research assistant Hailey Gabron, with input from the working group and other stakeholders. Raw result data was only available to the research team; full distribution data only available to the working group chair.

What should U-M be responsible for about AI?

In addition to disaggregating responses by faculty's self-reported field of expertise, as displayed in Figure F.1, responses were disaggregated by respondents' **position**, **time as faculty at the U-M**, as well as by **campus or Ann Arbor school or college**.

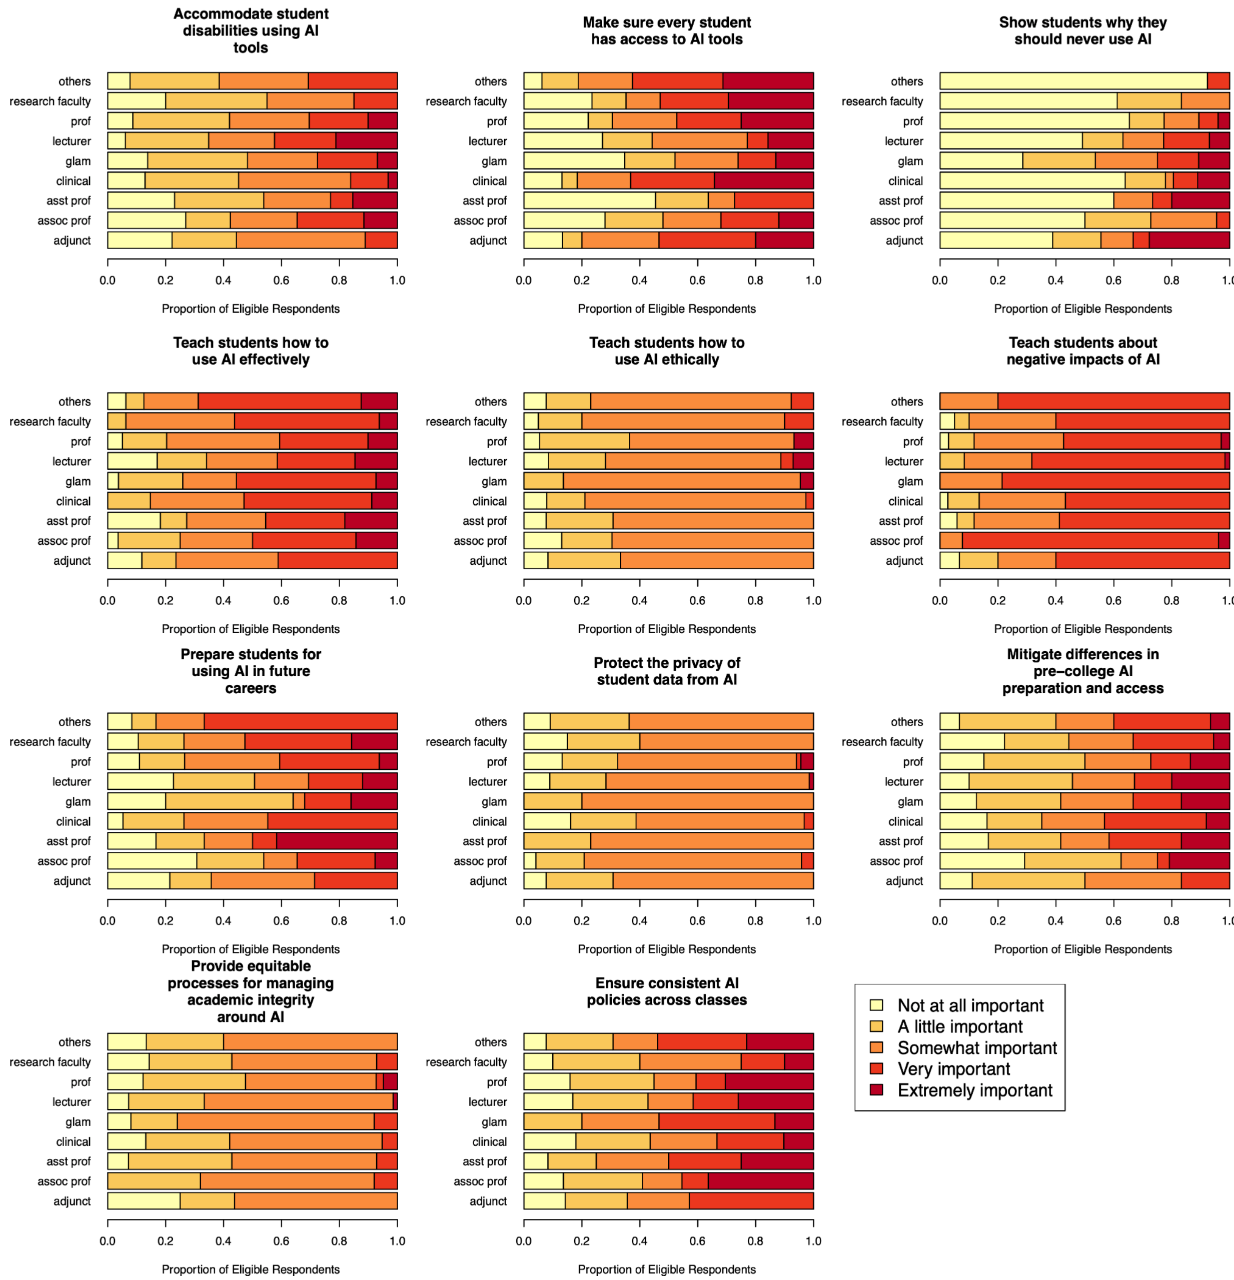

Figure F.1: Responses to “How important is it for the university to do each of the following?” disaggregated by faculty members’ *position*.

Similarly, faculty’s time at the university appears to make virtually no difference.

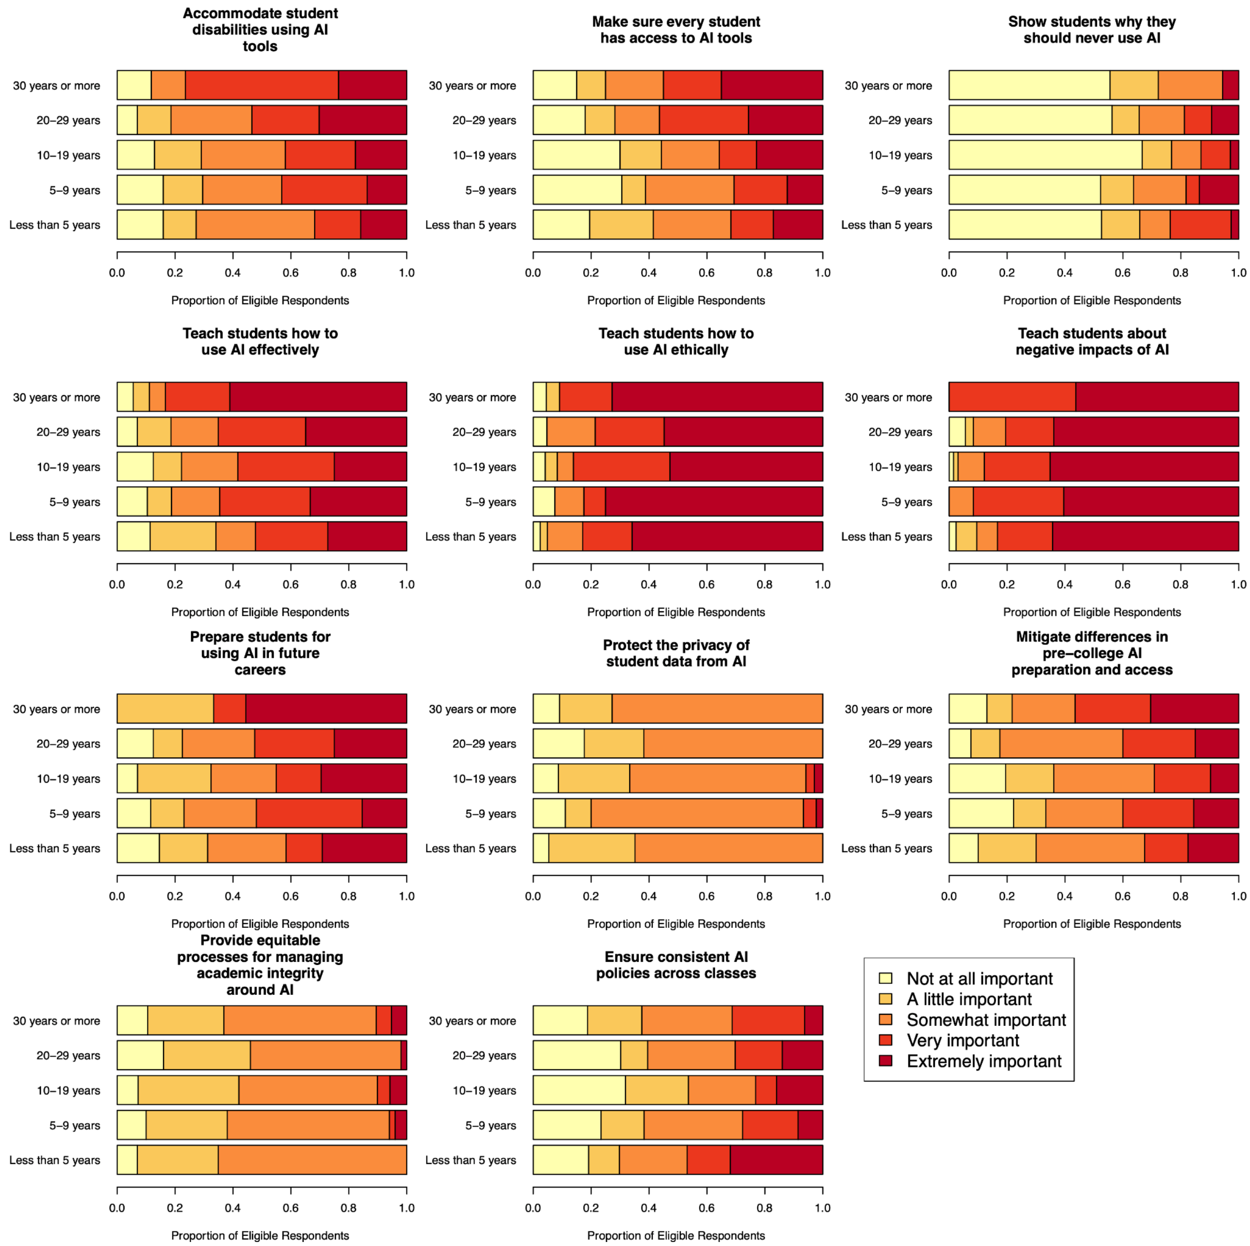

Figure F.2: Responses to "How important is it for the university to do each of the following?" disaggregated by faculty members' *time at U-M*.

Finally, faculty members' primary units seem to not make a difference, as the plot below suggests. This is particularly noteworthy, given the overrepresentation of LSA and underrepresentation of Michigan Medicine faculty in the survey.

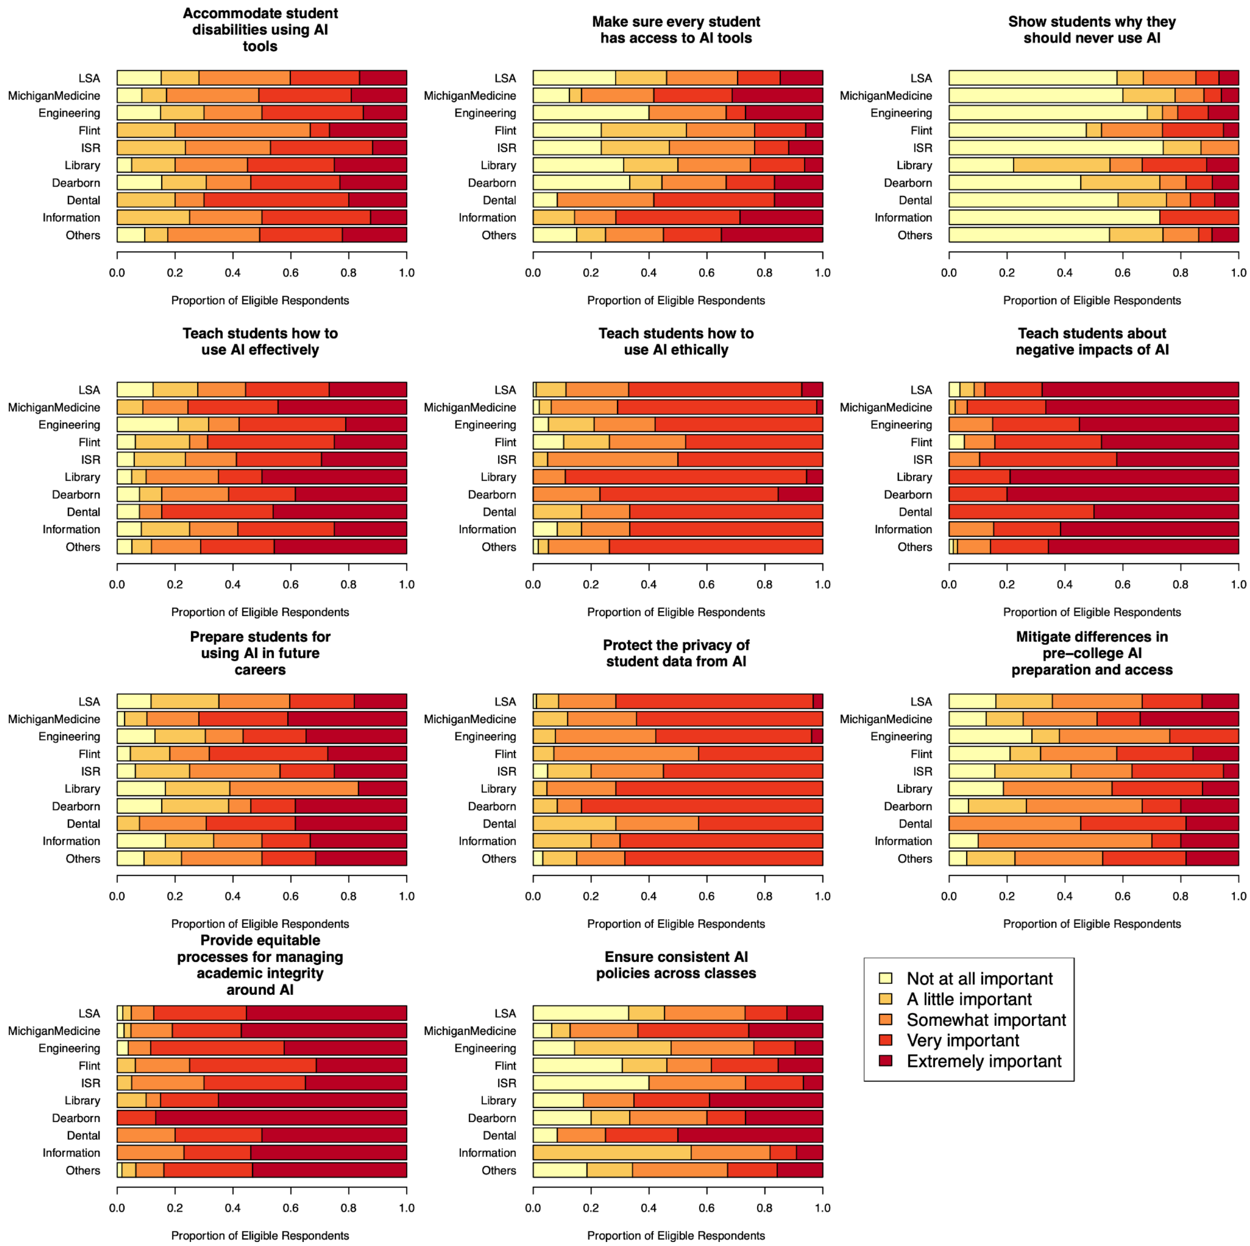

Figure F.3: Responses to "How important is it for the university to do each of the following?" disaggregated by faculty members' *campus, school, college, or administrative unit*.

Who sets policy?

The survey asked faculty their views on where they believe AI policy is set. Given that policies can be nested, they could choose multiple levels.

Three quarters of the faculty believe that policies are set by individual faculty members, about half also believe both the central administration and their school or college should play a role.

Who do you think sets policies around appropriate generative AI use for teaching and learning at the University of Michigan?

(Multiple choices possible; totals exceed 100%.)

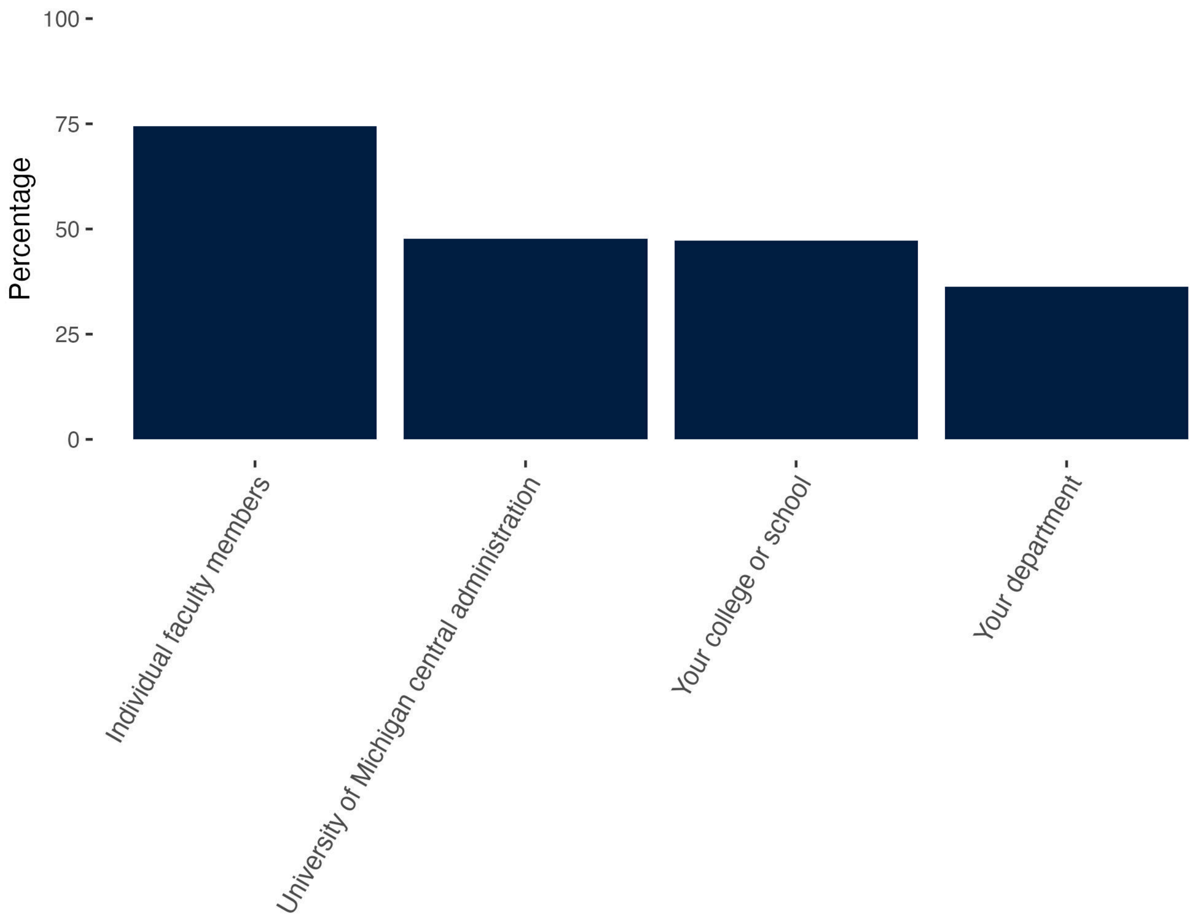

This roughly reflects faculty trust in these various institutional levels, with the exception that faculty trust their own department more than the higher levels of administration. For AI policies, the working group concurs: this report recommends that concrete AI frameworks for teaching and learning be set at the level that also oversees specific curricula.

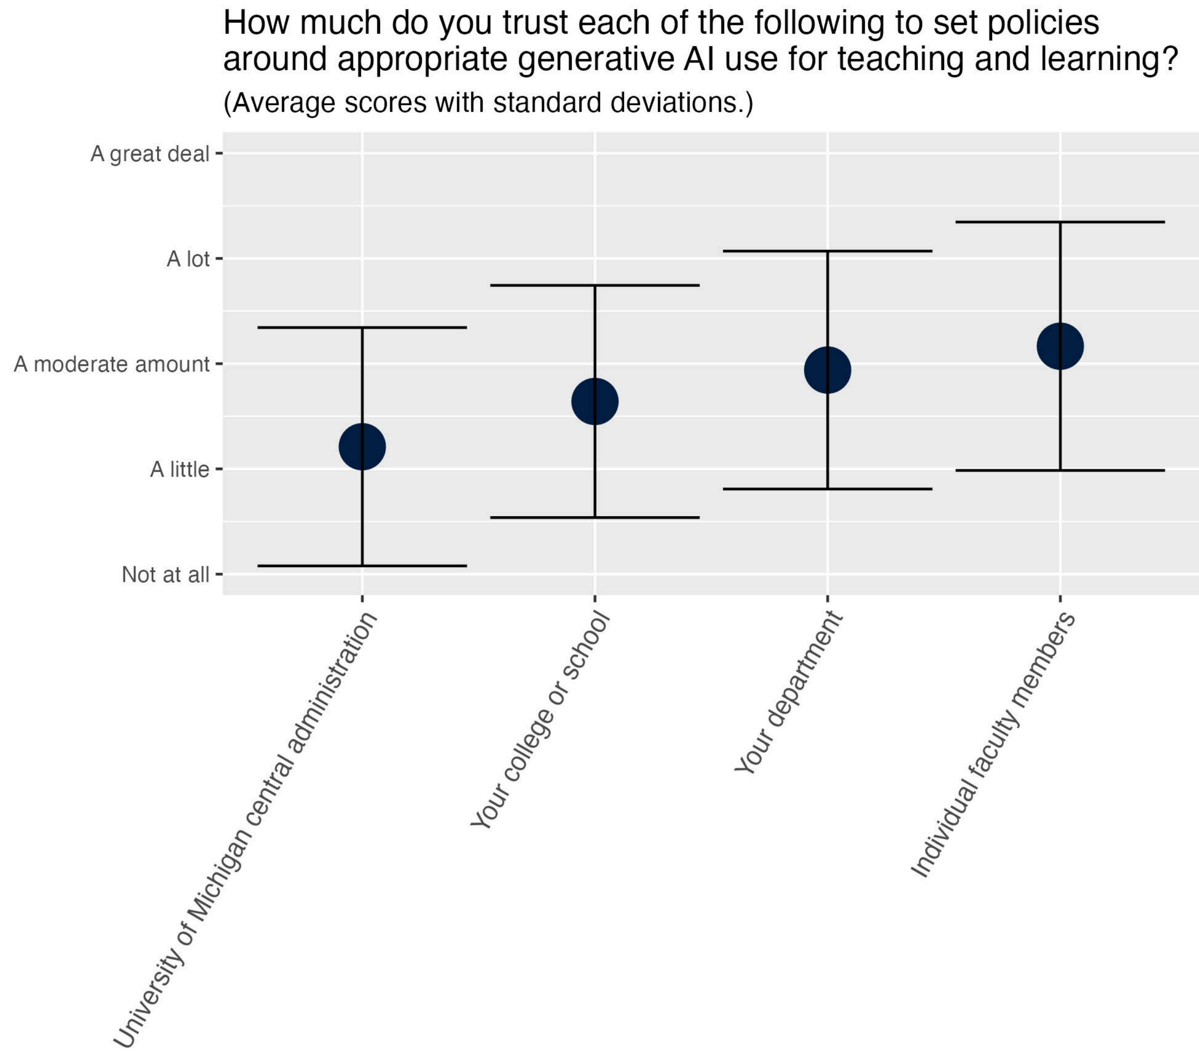

Open-ended comments

Below is a summary of the open-ended responses to the final open-ended question ("Is there anything else you would like to tell us?"). The open-ended questions have a negativity bias. Percentages indicated are from the **total** number of respondents, and represent a small minority. **Despite the small number, we include these here to represent what is clearly a strongly concerned minority.**

Results have been summarized by GPT 5.2.

Below are the main themes that emerged across the comments, grouped into categories that fit the substance of what respondents chose to emphasize. Percentages are *rough estimates* based on this set of pasted open-ended comments (respondents often expressed multiple themes, so totals exceed 100%). We include one illustrative quotation per category when helpful.

---

1) Strong anti-GenAI stance (seen as harmful / “not a panacea” / “existential threat”)

What respondents said: Many frame genAI as fundamentally misaligned with university aims (thinking, writing, scholarship), sometimes describing it as “snake oil,” “slop,” or inherently harmful to minds/learning. Several want U-M to “pull back” rather than integrate.
Illustrative quote: “AI is not the panacea that we were sold. Please move on.”
Estimated proportion: ~8-9%

---

2) Environmental/resource impacts (energy, water, data centers) as central objection

What respondents said: A large share of foreground environmental costs (water/energy, rare earths), sometimes pointing specifically to local data center concerns (e.g., Ypsilanti). Some argue environmental harms swamp other concerns like privacy/pedagogy.
Illustrative quote: “We simply can not talk about AI without also discussing its devastating impacts on our clean water and the environment.”
Estimated proportion: ~5%

---

3) Critique of “AI boosterism,” big-tech influence, and non-consultative/top-down rollout

What respondents said: Many distrust administration’s motives and messaging, describing an “all-in” stance, vendor capture, or “shoving AI down throats,” and express low confidence the survey will change decisions.
Illustrative quote: “The University has been shoving AI down the throats of faculty… and has spent undisclosed millions developing AI tools.”
Estimated proportion: ~6%

---

4) Academic integrity + inability to assess learning (cheating, “I can’t tell what’s student work”)

What respondents said: A major theme is that genAI undermines assessment validity and classroom trust. Respondents describe widespread outsourcing (homework, exams, discussion posts), students lying, and faculty time shifting from feedback to policing.
Illustrative quote: “GenAI has radically undermined my ability to assess student work… I spend more time trying to identify the use of AI rather than commenting and evaluating the work itself.”
Estimated proportion: ~10-12%

---

5) Perceived decline in core skills and learning dispositions (reading, writing, struggle, resilience)

What respondents said: Many argue AI accelerates “deskilling,” reduces persistence through difficulty, and erodes critical thinking and writing development—often described as the *point* of education. “Struggle” is repeatedly framed as essential for learning.
Illustrative quote: “My main concern… is that it takes away the *struggle*, which is a very important part of the learning process.”
Estimated proportion: ~9-11%

---

6) "Teach AI literacy / responsible use" (integration is inevitable; teach critical use, limits, ethics)

What respondents said: A substantial minority want proactive integration: teach students how to use AI thoughtfully, verify outputs, understand limitations/bias, and prepare for workplaces where AI is common. Some describe course designs where required AI use reveals its limits.
Illustrative quote: "The real task is to teach students how to use AI thoughtfully, responsibly, and effectively, while placing greater emphasis on critical thinking…"
Estimated proportion: ~5%

---

7) Context-dependence and need for nuance (discipline/course level/task differences; writing vs editing)

What respondents said: Many complain that "use AI" is too ambiguous and policies must vary by course, discipline, and task (especially distinguishing *editing* vs *writing*; ESL support; undergrad vs grad; different course formats).
Illustrative quote: "Some instructors allow GenAI for editing but not initial writing… questions require nuance."
Estimated proportion: ~6-7%

---

8) Calls for clearer policy frameworks + shared language (not one-size-fits-all mandates)

What respondents said: Respondents want clarity for students moving across courses: consistent "menu"/framework of policy levels, clearer academic integrity definitions, disclosure expectations, and guidance that also supports faculty who choose *not* to use AI.
Illustrative quote: "There does need to be policy of uniform prior information for each specific class about whether AI is allowed, and to what extent."
Estimated proportion: ~5%

---

9) Requested operational supports for "AI-resistant" assessment (testing space, TAC capacity, proctoring, tech)

What respondents said: Many ask for practical infrastructure: expanded Testing Accommodation Center (TAC) capacity, more in-person/secured testing facilities, better lock-down options, improved document/version history integration in Canvas, printing support for screen-free classrooms.
Illustrative quote: "We need a massively expanded TAC so that faculty can assign all blue-book exams…"
Estimated proportion: ~5%

---

10) Requested tools/support for AI access, privacy, and "approved" enterprise offerings (esp. for clinical/research use)

What respondents said: Another subset wants *better* institutional access to modern tools (Claude/Gemini/ChatGPT), stronger information assurance, and clear FERPA/privacy guidance—arguing that weak campus offerings push people to risky personal accounts.
Illustrative quote: "Faculty are effectively encouraged to explore these technologies with personal accounts—and that represents an institutional risk…"
Estimated proportion: ~3-4%

---

11) Workload, burnout, and morale impacts (time cost, extra policing, joy lost)

What respondents said: Faculty describe AI as making teaching "not fun," increasing workload (redesigning assessments, policing, accommodation logistics), and harming mental health. Several ask for time, grants, or compensation if expectations change.
Illustrative quote: "We need more time… It's like being asked to teach all new preps."
Estimated proportion: ~5%

Cross-cutting pattern worth noting

A recurring divide is **"embrace and teach responsible use" vs "protect learning by constraining/avoiding AI,"** with many commenters *not* purely in one camp—e.g., "AI is here to stay" but also "it's harming learning," coupled with requests for both literacy training *and* stricter assessment conditions.

Representativeness

Given an overall response rate of 8%, one may ask whether it accurately represents U-M faculty. Table F.1 below reports both the overall proportion of the unit of *all* faculty who received the survey invitation ("Percentage") and the unit's survey response percentage. Units with fewer than 20 respondents each were combined into the group "Other." Dearborn and Flint campuses were each combined due unit-size response rates.

The table shows that Michigan Medicine was significantly underrepresented in the results, while LSA is significantly overrepresented.

Table F.1: Survey response proportions and unit overall proportions.

Unit	Percentage	Response Percentage
Michigan Medicine	36.59	13.23
LSA	13.06	31.79
Engineering	7.22	8.47
Flint	4.54	7.08

Dentistry	4.49	3.25
Dearborn	4.45	4.87
Public Health	2.38	2.55
SMTD	1.91	2.44
Library	1.69	4.87
Ross	1.64	2.44
Information	1.38	2.9
Other *	0.21	0.16

The underrepresentation of Michigan Medicine may, in part, be explained by clinical faculty, who make up a large proportion of its faculty. Table F.2 below reports the eight largest faculty groups by faculty status. Some respondent groups have been omitted from the table; some, such as levels of Lecturers' Employee Organization (LEO) faculty or categories of clinical faculty have been grouped together. Clinical faculty is significantly underrepresented.

Also noteworthy are the overrepresentation of LEO faculty and full professors. Librarians are overrepresented, but due to their small proportion, not a potential source of bias in the overall results.

Table F.2: Faculty status overall proportions and corresponding response percentages.

Status	**Percentage**	**Response Percentage**
Clinical	38.35	14.39
LEO Lecturer	16.55	28.77
Professor	16.44	25.06
Research Faculty	8.19	5.92
Associate Professor	7.71	11.25
Assistant Professor	7.28	7.66
Adjunct	2.47	0.7
Librarian	1.72	4.76