

**Report of MIT's
Ad Hoc Committee on
AI Use in Teaching, Learning,
and Research Training**

August 13, 2026

Contents

1. Introduction	1
1.1. The Landscape	3
2. Guiding Principles	4
2.1. Be humble	4
2.2. Be bold	4
2.3. Put humanity front and center	5
2.4. Lean into learning	5
2.5. Teach with intentionality	6
2.6. No one size fits all	7
2.7. Augmentation not automation	7
2.8. Think beyond the classroom and the campus	8
3. Recommendations	9
3.1. Adapt educational processes for an AI-aware world	11
3.2. Center people, community, and the residential experience	18
3.3. Build processes, teams and tools for continuous reflection, iteration, and improvement	22
4. Conclusion	27
Appendices	29
A. Committee Process	29
B. Sample AI Policy for Syllabi	29
B.1. Unrestricted GenAI Use	29
B.2. Limited GenAI Use: Support Tool Only	30
B.3. Required GenAI Use	30
B.4. GenAI Use Strictly Prohibited	30
C. Survey Results	31
C.1. Spring 2026 AI Usage and Attitudes Survey	31
C.2. Spring 2026 Quality of Life Survey	34
C.3. Fall 2025 Tech Survey	37
D. Additional Resources	38

1. Introduction

This report is a call to action.

Through five intense months of meetings, research, and outreach across the MIT community, the Ad Hoc Committee on AI Use in Teaching, Learning, and Research Training sought to understand the role of generative AI in the life and educational mission of the Institute and recommend how to navigate its challenges and opportunities.

In January 2026, Chancellor Melissa Nobles, Provost Anantha Chandrakasan, and Faculty Chair Roger Levy charged us¹ specifically to:

- Assess current AI use at MIT
- Identify innovations in teaching and student assessment
- Propose an AI use policy.

However, what we learned as a group quickly convinced us that the Institute community, particularly the faculty, must tackle a set of deeper questions about the structure, meaning, and value of an MIT education in an era in which AI is one of several factors complicating the Institute's mission.

Our committee consisted of undergraduate and graduate students, faculty from every school, and staff from relevant units, including the MIT Libraries and the Teaching and Learning Lab. Though we brought to the assignment a broad range of experience and no fixed thesis, our brief but intense explorations led us to a strong shared view.

In this report, we:

- Highlight key aspects of the current educational landscape at MIT
- Share eight principles we relied on and that we hope will guide the Institute in the work ahead
- Recommend immediate and long-term actions for both instructors and the administration.

As an institution deeply identified with the birth of AI and known for its distinctively rigorous, hands-on education, designed to produce graduates unafraid of the world's hardest problems, MIT has a unique role to play in this moment. We believe it also has a responsibility to lead.

We hope our report can help the Institute lean into the spirit of Mind, Hand, and Heart as it continues to define and foster the highest-quality residential education – of humans, by humans, in support of human flourishing, and for the betterment of humankind.

Eric Klopfer, co-chair

Sam Madden, co-chair

On behalf of the Committee on AI Use in Teaching, Learning, and Research Training

¹<https://facultygovernance.mit.edu/committee/committee-ai-use-teaching-learning-and-research-training>

Committee Membership

Eric Klopfer, Co-chair; Professor, Comparative Media Studies/Writing Program

Sam Madden '99, MNG '99, Co-chair; Distinguished College of Computing Professor, Department of Electrical Engineering and Computer Science

Faez Ahmed, American Bureau of Shipping Career Development Professor, Department of Mechanical Engineering

Jacob Andreas, Associate Professor, Department of Electrical Engineering and Computer Science

Chris Bourg, Director of Libraries, MIT Libraries

Cynthia Breazeal SM '93, SCD '00, Professor, Media Arts and Sciences Program

Hope Dargan '19, MNG '23, PhD Candidate, Department of Electrical Engineering and Computer Science

Daniella DiPaola SM '21, PD '26, PhD '26, PhD Candidate, Media Arts and Sciences Program

Madalina Griza '26, Department of Electrical Engineering and Computer Science

Daniel Jackson SM '88, PhD '92, Thomas J. Perkins Professor, Department of Electrical Engineering and Computer Science

Graham Jones, Professor, Anthropology Program

Ankur Moitra M '09, PhD '11, Norbert Wiener Professor, Department of Mathematics

Caitlin Mueller '07, PhD '14, SM '14, Associate Professor, Department of Architecture

Janet Rankin PhD '89, Director, Teaching and Learning Lab

Eric So, Sloan Distinguished Professor of Management, Sloan School of Management

Marin Soljačić '96, PD '00, Cecil and Ida Green Professor of Physics, Department of Physics

Rodrigo Verdi, Nanyang Technological University Professor of Accounting, Sloan School of Management

Thomas Wu '26, Department of Electrical Engineering and Computer Science

Staff to the Committee

Jill Bassett, Chief of Staff, Office of the Chancellor

Martha Eddison, Special Assistant and Senior Communications Strategist to the President, Office of the President

Brittany Hutchinson, Executive Assistant, Office of the Chancellor

1.1. The Landscape

This report's recommendations reflect the following insights about the educational landscape at MIT.

Generative AI is everywhere already, spurring an assortment of views:

- MIT students use AI frequently and pervasively – with strongly mixed feelings, from curiosity, creative inspiration, and gratitude to resignation, concern, and anxiety.
- Instructors' attitudes range from enthusiastic exploration and growing reliance on AI to skepticism, suspicion, and "AI refusal" – and there's a widespread desire to share experiences, ideas and techniques.

While AI is allowing instructors to develop exciting new learning experiences and enabling students to learn and experiment in innovative ways, there have been a number of concerning effects on the life of the campus, including signs that AI is:

- Upending foundational elements of the MIT educational experience, especially for undergraduates, from the p-set, the take-home exam and UROPs to office hours and the study group
- Increasing isolation
- Undermining student mastery and confidence
- Eroding the "social contract" between instructors and students
- Making it much harder to assess student progress
- Challenging decades of distinctive MIT community norms and values about rigor, the creative friction required for learning, collaborative problem solving, and personal integrity.

AI is generating both immediate rapid changes and long-term tectonic disruptions – and MIT needs to respond:

- Students are confused and concerned about a lack of clarity, consistency and justification about the use of AI, within a given subject and across the curriculum
- Every subject taught at MIT will likely need to be reexamined and potentially revamped to make sure that how students are being taught, what they're learning, and how they're assessed are "AI-aware."
- AI is changing what students need to know and know how to do.
- Beyond its implications for specific subjects and disciplines, the advent of AI demands a broader, holistic reassessment of the nature, scope and purpose of higher education today.

Finally, AI presents substantial practical concerns for our community, from data privacy and confidentiality to questions around disparate access, bias, fairness, and accountability. AI's explosive growth and ascendance also raise important questions for society, from the environmental impact of AI data centers, to the use of intellectual property and training data, to the overall human impact of the technology and the industry – and MIT needs to engage with those questions too.

Other educational institutions are grappling with similar questions around AI and generating interesting ideas – but no one seems to have it all figured out.

Three notes on the words we use:

- *N.B. In this report, “Instructors” includes faculty and everyone else engaged in teaching at MIT.*
- *Most references to “students” apply to both undergraduate and graduate students, except in a few obvious places, as when we refer to the General Institute Requirements (GIRs) or to participants in the Undergraduate Research Opportunities Program (UIROP).*
- *“We,” “us,” and “our” sometimes refer to the members of the committee, and sometimes the whole of MIT. The difference should be clear from the context.*

2. Guiding Principles

We start by defining eight principles we relied on and that we hope will guide the Institute in the work ahead.

2.1. Be humble

Some technological innovations emerge gradually: As society and technology evolve in concert, mutual adaptation softens the impact. The computer – AI's precursor and key enabler – fits this pattern. Other innovations land more abruptly, becoming socially consequential before individuals and institutions have time to adapt.

Society tends to peg the “birth” of a new technology as the point when it becomes readily usable. By that measure, generative artificial intelligence was “born” with the release of ChatGPT in late 2022. Public engagement with generative AI is therefore less than four years old. In that time, it has amassed more than a billion users, and the companies selling AI technology have come to dominate the headlines, the stock market, and public consciousness.

In other words, AI is progressing across almost every domain and on a timescale too compressed for society to properly observe and analyze its impacts and then gradually adapt.

This suggests our first guiding principle: We offer our proposals in a spirit of humility. Course corrections – perhaps even major ones – will be inevitable as the technology continues its relentless evolution and the Institute experiments and learns.

2.2. Be bold

Yet uncertainty can't be an excuse for inaction. This is not a moment for patches and duct tape. The challenges AI presents in teaching and learning call for a bold strategic response – everywhere, and especially at MIT. With our Social and Ethical Responsibilities of Computing program² completing its seventh year, we are uniquely positioned to find ways to employ this new technology for the benefit of society, and for our students in particular.

²<https://computing.mit.edu/cross-cutting/social-and-ethical-responsibilities-of-computing/>

AI also presents extraordinary opportunities, from unprecedented possibilities for individualized tutoring and coaching to a dramatic acceleration and revamping of research in many disciplines. Seizing these opportunities deserves and demands boldness too.

Bold thinking is especially important because our students will go on to help shape the intellectual, ethical, and technical direction of our society – and soon. We owe them a deep engagement in rich and constructive uses of AI, and a sophisticated understanding of its potential and its drawbacks. Their MIT experience should prepare them with the wisdom to help determine how and where AI is used for the betterment of society and the world at large.

2.3. Put humanity front and center

Facing a technology that already has such immense capabilities – built on, modeled after, and in many ways now exceeding human powers – the Institute's fundamental challenge and most important goal must be to nurture and protect our shared humanity, and to value the MIT community, its members, and their flourishing above all.

In a listening session with instructors, we learned that some were considering using AI agents as research assistants instead of hiring undergraduates as UROPs. One can see the case for speed and efficient use of resources (especially now, when research resources are so constrained). But if those criteria come to dominate our decisions, we all have to ask, "What is it that we are here together to do?"

As a community, we need to keep in mind that although research is central to MIT's mission, it's more than an end in itself; on our campus, research is also an apprenticeship, a means of training the next generation of researchers to continue our work and drive our fields of inquiry forward. Using research as an opportunity for learning-by-doing may produce seeming "inefficiencies," but that's a feature, not a bug.

A focus on our humanity should also alert us to the fact that common current remedies for the problems AI creates can risk damaging the relationships between and among teachers and students. For instance, many instructors recounted that having to "police" unauthorized AI use was harming their connection to students (a dynamic made worse by the fact that, as we learned, AI detection software is quite unreliable). For their part, students fear being wrongly accused of AI cheating and are increasingly frustrated by instructors' use of AI in areas that demand a human touch, such as grading, assignment creating, and feedback. Such an underground river of mutual suspicion is no foundation for a healthy classroom.

2.4. Lean into learning

Many problems and assignments used in MIT classes to reinforce learning and assess student progress can already be accomplished by AI. Unrestricted AI use by students can make some traditional assessments less reliable as indicators of individual learning, potentially weakening confidence in grades and credentials. For educators, this radical shift in norms and expectations can feel profoundly disorienting.

But as a community, what should worry us most is that many uses of AI deprive students of the opportunity to learn.

The threat AI poses to familiar ways of teaching and testing may be a blessing in disguise – because the changes are too sudden and severe to ignore. As the faculty on our committee can attest, for at

least two decades, educators have lamented the strain on teaching and learning from fragmented attention, ubiquitous devices, a narrow preoccupation with grades rather than learning, and a weakening sense that students and teachers are joined in a common intellectual project.

As an institution that prides itself on the power of its distinctive educational recipe, it is up to MIT to seize the opportunity of this moment: to make sure that the undeniable changes imposed by AI become a tipping point, forcing us to deal decisively with the forces eroding our shared educational mission.

Leaning into learning means creating a new “social contract” between teachers and students. All of us who teach at MIT will need to be prepared to help students understand both that the *process* of education is necessarily a productive struggle, and that the most important *product* of their education is not a GPA or a diploma but *themselves*: their personal growth and intellectual maturity and the development of their own imagination, insight, and judgment.

Instilling these attitudes needs to become a central task for every educator, so that our students know not only *what* they should learn, but also *how* they should learn and why learning matters. We need to help them develop metacognitive abilities to think about thinking, to engage in reflective practices, and to enhance their sense of personal agency. This will require both dedication and fresh preparation on the part of instructors.

2.5. Teach with intentionality

Most classes taught at MIT will require careful review and, in many cases, substantial adaptation to survive an encounter with AI without serious disruption. AI will undermine the pedagogical assumptions that guided the choice of content, the methods of instruction, and the ways of assessing student performance.

In short, it's time to reflect deeply and extensively on what we teach and how we teach it. Future class structures, assessments, and policies must be built with high intention and a clear sense of purpose, not merely tweaked in reaction to the immediate realities of AI. This does not mean, however, that every course must change in the same way.

The time-tested technique known as “backward design³” offers a useful framework here. For instance, rather than beginning by asking whether AI should be allowed or prohibited in a specific subject, educators would begin by defining the purpose of the learning experience itself: what students should come to know, be able to do, and learn to value. Assignments and assessments can then be designed to measure and encourage those outcomes, using AI if it's helpful but not if it isn't.

The instructor time and effort required to redesign, reassess, or intentionally preserve existing structures is likely to be significant. But that process of pedagogical reflection and reasoning is valuable in itself. Intentional teaching promotes intentional learning; when instructors make clear why AI is permitted, limited, or required, students are more likely to understand the learning that's being protected or developed.

In this way, backward design can help MIT preserve the purpose of its education while adapting honestly to the realities of AI. (Encouragingly, MIT's Teaching and Learning Lab is well-equipped to help instructors master this tool.)

³<https://tll.mit.edu/teaching-resources/course-design/where-to-start-backward-design/>

2.6. No one size fits all

No single approach to AI in education will serve MIT well. A poetry seminar, a mathematical proof course, and an architectural design lab logically entail different relationships to AI. In some, using AI to generate creative work or solve a problem defeats the purpose of the exercise; in others, AI fluency is already part of professional practice, and avoiding it would leave students shortchanged.

Similarly, a first-year student building foundational skills and judgment stands in a different relationship to AI than doctoral candidates accelerating literature review in a field they deeply know. In short, a uniform Institute-wide rule would inevitably be too permissive for some contexts and too restrictive for others.

This is not to say that every instructor should have to invent AI policy from scratch. MIT should provide a shared framework – including a common policy menu, disclosure expectations, and accountability standards – within which departments and instructors can make choices. Departments are a natural unit for this work, since students experience a major as a coherent progression and would benefit from consistency about when AI is welcomed, when it is restricted, and why.

2.7. Augmentation not automation

As AI-enabled technologies become more capable, it will be possible – and tempting – to offload more “thinking” tasks to them; not surprisingly, students told us the temptation was greatest when they feared they would miss a deadline. There are early signals, however, that overreliance on chatbots can have a range of significant negative consequences – diminishing critical thinking, weakening memory, eroding confidence, and undermining mastery.

Getting the right answer from a chatbot can create the illusion of learning – but it can also trigger “cognitive surrender”⁴, where students fall back on AI at the first hint of struggle. In a recent paper⁵, three noted MIT economists – Daron Acemoglu, David Autor and Simon Johnson – argue for what they call “pro-worker AI.” They make the case that instead of designing AI systems to replace workers, businesses should develop AI that allows human beings “to be more effective at their existing tasks, tackle new tasks, and acquire new expertise.”

In the same spirit, collectively, we need to shape MIT’s use of AI in education to be “pro-learner” – not to replace the “hard fun” of thinking, learning, creating, and problem solving, but to expand what’s possible for students to think about, learn, and solve. In short, AI should be used to augment and enhance curiosity, creativity, and learning, not automate them.

Importantly, this issue is social as well as psychological. Just as automation threatens job loss, deskilling, and the increasing isolation of workers, so too the reliance on chatbots in education displaces peer-to-peer connections like study groups and decreases interactions between students and instructors. If we want augmentation to carry the day instead, we need to see it as more than an abstract, “hoped-for” ideal; it’s a concrete practice that our pedagogies can help promote.

This leads to our final principle.

⁴https://osf.io/preprints/psyarxiv/yk25n_v1

⁵<https://www.brookings.edu/articles/building-pro-worker-ai/>

2.8. Think beyond the classroom and the campus

AI should be integrated into MIT education in ways that enhance, rather than diminish, students' capacities for growth as whole human beings. In *The Culture of Education*⁶, Jerome Bruner, a leading psychologist of education and a pioneering figure in cognitive science, argued that learning is not simply a technical matter of transmitting content from teacher to student, but a process through which students learn how to interpret the world, find their place in society, and develop a sense of themselves as capable persons.

If learning is (mis)understood primarily as content delivery, then highly responsive AI systems may seem capable of replacing teachers, mentors, peers, collaborators, and the interpersonal dimensions of learning altogether. But while this attitude may be increasingly common, it's shortsighted. Education is a cultural practice through which students learn to make meaning, exercise judgment, form identities, and participate responsibly in community.

Much of what our students gain from MIT is never spelled out in a syllabus or an assignment; it's what they learn from living and working on our campus in each other's company – the tacit expectations, habits, relationships, and values that inform how they learn to solve problems, exercise judgment, persevere through difficulty, and become members of an intellectual community. "Residential education" is powerful in part because it happens everywhere: in residence halls, living groups, sports teams, arts groups, clubs, and so on.

A kind of social contract

During our discussions, we came to think of the alignment among these elements as a kind of social contract: a shared understanding among students, instructors, and the Institute about the purposes of education, and the rights we enjoy and responsibilities we owe as members of an intellectual community. Effective institutional responses to issues under consideration must therefore attend not only to policies and rules governing AI use, but also to the norms and culture that give those policies meaning.

A fundamental danger, as we've discussed, is that AI can allow students to bypass learning. Equally concerning is that students may internalize a transactional model in which assignments are outputs, teachers are evaluators, peers are optional, and knowledge (or an MIT degree) is an optimizable commodity to be acquired or produced as efficiently as possible.

Such a mental model will not remain confined to the classroom. It will shape how students come to understand work, collaboration, and social responsibility, and they will carry that mindset with them out into the world. However sociable or responsive AI systems become, they cannot substitute for the relationships and practices necessary to grow and mature as a human being.

MIT should therefore approach AI not simply as a tool for improving educational efficiency, but as a force that must be governed to support the broader purposes of education. Students should learn to use AI productively and fluently, but also in ways that deepen their capacity to participate well in their classrooms, laboratories, workplaces, families, and communities.

The goal is not to shield students from AI, nor to preserve older educational forms for their own sake. It is to ensure that AI use supports the development of people who can think critically, act with initiative, work productively with others, and understand the consequences of their choices in a world shared with 8.3 billion other human beings.

⁶<https://www.hup.harvard.edu/books/9780674179530>

3. Recommendations

This section details the changes we see as necessary for MIT to prepare our students for a new world. Based on the Institute's longstanding values, our recommendations are not a checklist of individual initiatives that can be implemented one at a time, bit by bit, but rather a set of substantive changes that must be undertaken in concert.

We recognize that serious change takes time. Given the impacts already affecting our community, however, the required changes should be implemented on two timescales: those that happen immediately and those that begin immediately, but require further study and planning.

Artificial intelligence – in the form of LLMs and other generative AI technologies – presents MIT with profound challenges and intriguing opportunities.

Already these technologies can produce credible solutions and provide reasonable responses to almost any written assignment in our undergraduate curriculum, including essays, math and science problems, proofs, and coding assignments, and their power will only grow.

Concerning impacts

Because many students are choosing or feel pressure to shift to learning and problem-solving with AI, in less than three years these technologies have driven major shifts in campus culture, including decreased attendance at office hours, reduced participation in online discussions, and, as we heard anecdotally, a drop in in-person study groups in dorms, libraries, and other study spaces. These issues have presented themselves suddenly and dramatically, creating a clear sense of urgency. They also land at a time when higher education is facing other challenges, and MIT itself is considering broad curricular changes emerging from the findings and recommendations of the Taskforce on the Undergraduate Academic Program (TFUAP)⁷.

Intriguing opportunities

While these challenges are pressing, AI also offers exciting opportunities for learning and for creating. Many instructors told us that AI helps them develop customized, interactive learning tools that allow students to explore subject content with more depth and for instructors to create learning experiences for their students that are new or newly tailored to each student. With a little guidance, even instructors who are not proficient in software development can customize AI agents to support a subject or research project. Students can create large-scale software projects with the limited timeframe of classes that would not have been remotely possible before. They can also use AI to analyze data, conduct research, and build tools to amplify their expertise and pursue projects that benefit them and society at large. The potential of these technologies to augment work across campus is immense.

The recommendations we offer below aim to help the MIT community navigate AI's challenges and seize its opportunities, while reflecting the principles laid out in Section 2. For instance, **Humility** means that we must create administrative processes for continuous evaluation and revision rather than assuming any change we make today will be adequate in the future. **Boldness** requires us not simply to try to patch the existing system to limit or counteract the impacts of AI on our students and systems but instead to redesign learning experiences, assessments, and curricula in ways that help redefine the future of education. **Putting humanity front and center** means that we will seek uses of AI that strengthen rather than weaken the value of an MIT residential education. **Leaning into learning** means teaching with intentionality, revisiting what we need students to learn, and aligning assessments with those desired outcomes, to preserve the productive struggle

⁷<https://gue.mit.edu/tfuap/>

essential to a learning-by-doing education. **No one-size-fits-all** means that we should develop tailored frameworks for adapting to AI rather than uniform AI rules for all students, instructors, or departments. **Augmentation over automation** means that we prepare students to use AI fluently and in ways that preserve agency, judgment, integrity, and human connection. **Thinking beyond the classroom and the campus** requires that we anticipate the skills that our students will need for success in work and life, as community members, future leaders, and creators of the next generation of AI technology.

Three broad areas of work

We believe the circumstances call for a substantial reshaping of the student experience. While the life of an MIT student will no doubt look different in 10 or 15 years, as a community we have a responsibility to make sure that it remains immediately recognizable as an MIT experience. Paradoxically, preserving what's most important will require deliberate change.

We group our recommendations into three sections, briefly summarized below. Broadly speaking, we recommend that, through the combined efforts of the faculty and the administration, MIT:

Adapt educational processes for an AI-aware world (Section 3.1)

- Campus-wide and as soon as possible, review what our students need to learn. Led by instructors and departments, this “AI-aware” process will inform the design of classroom activities that are more social and experiential, and in tune with innovative assignments and assessments that provide feedback on those key student outcomes. This will also require an investigation into the nature of grades and grading.
- Facilitate this process by providing instructors with dedicated resources (for example, additional TAs) and by adapting physical spaces for experimentation, community-building, and AI-free interaction and assessment.
- Create “communities of practice” at MIT that allow members to share challenges and findings, as well as tools and techniques to improve teaching and learning.

Center people, community, and the residential experience (Section 3.2)

- Reinvigorate and reclaim the residential experience based on shared experiences, transparency, and a focus on in-person activities.
- Thoughtfully integrate AI literacy and responsible and ethical AI use as a foundation and throughline for student work, attending closely to how it connects to specific disciplines. This will demand technical proficiency, of course, but it hinges on adopting a human-centered approach to AI as an augmentative tool.

Build processes, teams and tools to enable continuous reflection, iteration and improvement (Section 3.3)

- Establish roles, processes, and ongoing committees to address these challenges
- Collect feedback and data, stay apprised of research and best practices, connect with peers and learn from them
- Provide technical and educational support to enable all of this.

In addition, we identify a set of concerns about the wider impact of AI on society, jobs, the environment, and more, that also bear on our on-campus AI policy.

3.1. Adapt educational processes for an AI-aware world

Generative AI affects what our students need to know and be able to do (our learning outcomes), how our students achieve those outcomes, and how we evaluate and assess what they have learned. Because AI's ability to competently complete MIT-level assignments makes it difficult to assess student progress based on out-of-class work, instructors urgently feel the need for new assessment strategies.

While we make several recommendations about assessment, we urge instructors to do more than simply try to "AI-proof" their classes. In particular, the ways that AI changes the social aspects of learning – such as how students interact with TAs and engage in office hours and study groups – present an even deeper challenge.

3.1.1. Revisit course goals

Before considering how to construct AI-aware assessments, instructors should reconsider their goals for student learning⁸ in every subject they teach. What should students know or be able to do by the end of the course?

The learning goals for a given subject may or may not be affected by the availability of AI, but they should be "AI-aware"; that is, they should recognize that AI exists in the world, that instructors may direct or permit it to be used in or out of class in structured ways – and that students may also seek to use it without permission.

We heard that the greatest concern for many instructors is: When AI makes it possible to offload the cognitive work of learning, how can we assess what students actually know and understand?

In the era of AI, some traditional learning goals may merit rethinking; for example, do the majority of our students need to be able to write complex programs by hand? Individual instructors and departments will need to assess whether longstanding goals are still important and, if so, develop AI-resilient ways to evaluate them.

While learning goals that involve lab, project, or performance-centered work may not require much adjustment, generative AI may offer novel ways to help students reflect on, practice, or prepare to perform. For instance, an instructor reported to us that, for students learning to serve as mediators, providing personalized, course-specific AI coaches eliminated the awkwardness of practicing public speaking in front of others, which substantially increased students' willingness to practice, which in turn increased their skills.

AI may also enable instructors to devise learning goals that were previously impossible, such as understanding or working in new ways with very complex texts, engineering artifacts, or large software systems.

⁸<https://tll.mit.edu/teaching-resources/course-design/where-to-start-backward-design/>

3.1.2. Ensure durable learning through new course policies, structures, and forms of assessment

MIT's curriculum is famously demanding. On top of that, many students pack their schedules with high unit loads and multiple extra curricular activities. These competing demands on their time drive students to prioritize efficiency – and nothing could be more efficient than automating work through AI. But if students give in to that tempting option, they cheat themselves of the cognitive friction and productive struggle necessary for actual learning.

AI has both created and revealed a mismatch between established learning objectives and familiar forms of assessment. At MIT, we have long used problem sets, long-form written assignments, take-home exams, and projects completed outside of class to allow students to solidify their knowledge through practice and demonstrate what they've learned. AI saps these tools of their value for both teaching and assessment.

Instructors already sense it: In the face of AI, they have to change how they evaluate students. Already, many are increasing the weight of exams in their grading or asking students to write or code during class time. However, these tactical solutions come with a cost: For instance, overemphasizing in-class evaluations means reducing students' incentive to invest themselves in the difficult, time-intensive p-sets and projects it takes to build the ladder to mastery. By definition, shifting assessments to time-limited class periods reduces how much thought and deliberation students can put in. If we want students to care about and know how to create and recognize worthy work – work of scope, rigor, creativity and thoughtfulness – quick, high-stakes evaluations embody the opposite of the signal we want to convey to them right now.

Instead, instructors need to revisit what they really want students to know and devise assessments that foster, or even include, the kind of productive struggle that builds durable understanding and mastery.

We urge instructors to consider forms of assessment that are less vulnerable to AI, and more valuable for learning, such as oral exams, semester portfolios, and out-of-class assignments paired with in-class conversations. This likely means that resources such as TAs and class time will become more central to evaluation and also raises the question of whether we should find ways to limit class sizes.

As described in Section 3.3 below, we also hope that the Institute will create efficient ways for instructors to share new assessment ideas and best practices that may emerge.

3.1.3. Emphasize experiential and project-based learning

To match the shift towards assessment methods that aren't vulnerable to AI, instructors need to increase the role of experiential and project-based learning. To encourage this kind of creative teaching, MIT should support the development of teaching skills and practices for all instructors, and recognize contributions in this domain.

Ironically, in some cases AI itself can help open the door to experiential opportunities. For instance, since AI excels at tasks like coding and some types of design, instructors can now assign projects that are much more ambitious: in MIT's capstone level software engineering class, projects that students could reasonably undertake in a semester often had to be limited in scope and unrealistic in key respects. Today, it's entirely reasonable to expect students to use AI coding tools to build near production-quality software artifacts in a single academic term. This in turn can make room for deeper experiential learning: students can explore how different designs work in the real world,

reflect on cases where certain designs outperform others, and explore how the systems they build work in realistic settings.

Opportunities extend well beyond computer science: Architecture students are using AI to experiment with new ways to visualize and rapidly test their ideas, beyond what's possible with traditional representational skills. This expands the range of creative projects they can undertake (though it only increases the importance of building fundamental conceptual and technical skills, judgment and social reasoning, to ensure students remain in charge of their ideas.)

In a recent Wall Street Journal interview⁹, Harvard economist Rafella Sadun observed that workers who will succeed in the AI era are those “who can creatively think about ways in which AI can generate value, quickly translate ideas into action, and who are ready to revise their intuition through experimentation and validation. Tinkerers and doers thrive in this environment” – an obvious match with MIT's hands-on ethos and focus on societal impact. As AI enables our students to tackle more sophisticated and challenging tasks, one way to live up to our *Mens et Manus* motto will be to expand the scope and real-world application of the projects and endeavors we assign.

Finding new opportunities for collaborative, in-class projects provides significant opportunities for growth beyond mastering the particular content. Properly structured and supported, such projects can build students' strength in collaboration, communication, problem-solving, and emotional intelligence. These durable human skills remain among those that employers prize the most. We would do well to provide more opportunities for our students to build them.

AI can also enhance these experiences, offering students new ways to analyze data, conduct research, and build tools to amplify their expertise and pursue projects that benefit them and society at large. The potential of this technology to augment work across campus is immense.

3.1.4. Build structured in-person social learning into subjects

Because AI tools are disrupting longstanding patterns of social learning at MIT, every subject should include a regular in-person social component (not just sitting in lecture or recitation and quietly taking notes). Instructors should intentionally structure such interactions to achieve desired learning objectives and maintain quality, even in large classes.

Examples of structured social learning include:

- Group projects with weekly staff check-ins and deliverables that assess both individual and collaborative contributions
- Group problem-solving sessions guided by a lab assistant or TA
- Feedback discussions structured around a class rubric
- In-class discussions facilitated by a trained TA or professor, where individual participation is graded.

To get the most value from in-person interactions like these, instructors should clearly explain on the first day of class why they are important for individual learning and building community. Courses that include extensive group projects should offer practical guidance and structures for how to successfully work together.

⁹Chen, T.-P., & Lahart, J. (2026, June 9). Economists weigh in on the future of work and AI. The Wall Street Journal. <https://www.wsj.com/tech/ai/economists-weigh-in-on-the-future-of-work-and-ai-f59311e9>

Although the presence of AI is disrupting social learning on campus, AI tools can also offer real benefits for student learning, including timely help, individualized feedback, accessibility, and support for students who might otherwise be stuck, as well as just-in-time ways to help groups collaborate.

Unfortunately, these benefits become counterproductive if they allow students to bypass the human settings where they would learn how to work with others, communicate their ideas, receive criticism constructively, build confidence, develop judgment, and act as members of a community.

MIT should therefore treat AI not only as a pedagogical or technical challenge, but as a critical opportunity to renew residential education around human presence, shared work, and meaningful mentorship.

3.1.5. Preserve and expand out-of-class research and career experiences

With its launch in 1969, MIT's Undergraduate Research Opportunities Program¹⁰ (UROP) broke new ground by creating a systematic way to give undergraduates hands-on experience with frontline faculty-led research. Today, UROP is a signature part of the MIT experience, directly engaging 93% of undergraduates and 58% of faculty.

UROP and other out-of-class research and learning experiences offer students obvious practical benefits such as academic credit and paid work. **But the broader benefits – including personal development, connections with mentors and peers, and career exploration – are even more important. UROP makes an eloquent case for the purpose and value of residential college education.**

MIT should preserve, renew and ideally expand its commitment to UROP and other learning experiences that emphasize mentorship, collaboration, and learning by doing; find ways to make them more accessible; and consider extending them beyond the lab.

We also encourage the Institute to explore related efforts to enhance the value of the MIT residential experience, such as establishing co-op programs (which the Institute is actively exploring¹¹), allowing UROP-like experiences to satisfy degree requirements, or expanding industrial/co-op based graduate research programs. This should include working with organizations on campus that already support experiences like these, including the PKG Center for Social Impact and MISTI.

UROPs, research assistantships (RAs), and other forms of student participation can give students entry into new intellectual communities: joining a lab, learning from graduate students and postdocs, collaborating with peers at different career stages, presenting work in group meetings, contributing to papers, traveling to conferences, and gradually coming to see themselves as members of a field. Students learn not only methods and technical competencies, but also how research questions are formed, how judgment develops, how mistakes are constructively interpreted, how credit is shared, how disagreement is handled, and how knowledge is collectively produced.

One note about UROP in particular: While undergraduates can provide faculty with useful research labor, that is not the point of the program. It exists to educate. This is why the possibility that faculty may find it advantageous to replace novice researchers with AI agents warrants special concern.

¹⁰<https://urop.mit.edu>

¹¹<https://orgchart.mit.edu/letters/announcing-mit-co-op-planning-committee>

As AI systems become a cheaper or more efficient replacement for UROPs or RAs, students could lose access to the relationships, practices, and shared forms of work through which belonging, confidence, judgment, and professional identity are formed. The danger is not only that students would have fewer opportunities to conduct research, but fewer opportunities to become participants in research communities.

3.1.6. Reconsider grades and incentives

We do not advocate limiting the grades instructors assign through grade rationing (e.g., imposing a limit on the number of A's that can be given). The committee felt that this would be counter-productive to campus culture and to the Institute's efforts to navigate the challenges of AI. Due to the intense societal focus on credentials and grades, students who are set on maximizing their GPA have a strong incentive to use whatever means they feel are most effective to achieve that goal. Rationing top grades would intensify the temptation to cut corners on actual learning by increasing reliance on AI.

As instructors explore new forms of assessment and experiential learning activities, MIT should take this opportunity to consider what role grades play in our overall system, and if the current approach could be improved. For instance, if instructors mainly respond by putting more weight on in-class exams, there's a risk of narrowing what the MIT credential has long signaled: that MIT students are capable of difficult, independent, and thought-intensive problem solving, not just acing exams on paper.

The committee encourages MIT to explore alternative systems of grading and assessment, including those that work in other countries (such as the UK system, which uses percentages to express relative mastery against an expert standard), and newer paradigms such as competency-based and mastery-based assessments¹². As one thought experiment, the committee discussed the idea that if MIT did not have grades, many of the incentives around AI cheating would disappear. Note that MIT already refrains from the common practice of awarding *summa/magna/cum-laude* diplomas, because we believe that earning an MIT degree stands as distinction enough.

Grades are far from the only way for students to demonstrate mastery or gauge their own progress; real-time feedback during an oral exam or presentation, or written comments reflecting on a project submission, likely give students more useful information about their mastery of concepts than they could glean from any grade.

With an increased focus on projects and experiential learning, instructors should consider whether grades are the only, or even the best, mechanism for signaling mastery in all contexts. Alternatives could include elevating the standing of portfolios of projects as a showcase of student achievement in every field. Already, many employers are less focused on applicants' grades than on their performance on internal assessments, such as the ability to answer difficult interview questions or demonstrate problem solving and mastery through custom exercises.

3.1.7. Expand in-person spaces for labs and in-person evaluation

In several ways, our recommendations call for more face-to-face activities, from hands-on learning and collaborative lab projects to in-person assessments. All of them require physical space.

Several instructors expressed a desire for collaborative spaces where students can complete extended assignments or exams. These might be fully analog spaces or spaces with AI-free (or AI-limited)

¹²https://en.wikipedia.org/wiki/Competency-based_learning

computers. Many departments have lab spaces that could support collaborative activities but lack resources to staff them full time. (In fields that use expensive and possibly hazardous equipment and materials, self-service labs are often not possible.)

Collaborative hands-on spaces are useful for teaching, learning, and community building across the disciplines, not just in the traditional lab sciences. We urge MIT to invest in creating and staffing such spaces.

3.1.8. Provide AI use policies, with justification

Students report that the AI guidance they receive from instructors is often confusing and unclear. What's more, the rules and the way they're conveyed vary widely from instructor to instructor.

While we recommend against imposing a one-size-fits-all policy on AI use, students are anxious for clarity about AI use: in any given course, they want to feel sure about when, where, and why AI is prohibited, allowed, or required.

Therefore, instructors, and perhaps departments, should make sure that every MIT subject has a clear policy about the use of generative AI, posted prominently in the syllabus and on the course website.

To make this workable, as soon as possible MIT needs to develop a clear and consistent menu of guidelines for instructors and departments to choose from and adapt as necessary. Departments will likely want to take a coordinated approach so that guidelines and rationales are well understood and largely consistent across a given major.

We encourage instructors to use a standardized format to communicate this information, so students can easily understand what is and is not permitted, course by course. Appendix B below provides an example policy.

AI policies should include a rationale. While many courses already have an AI use policy, we recommend that all instructors make sure that their policies include a clear rationale, tied explicitly to the learning goals of the given course, for *why* AI must, may, or may not be used.

For example, if an instructor wants to ban the use of generative AI tools instead of simply declaring that AI use is a form of cheating, it's more effective to explain how generative AI tools shortcut students' ability to learn the fundamentals of the course or to practice solving the kind of problems they will encounter in exams, future classes, and the real world. This approach is more likely to foster self-awareness and a healthy culture around AI use.

Similarly, instructors who want to encourage AI use in a particular assignment should make it clear why using AI is important to the educational experience. For assignments that feature AI, instructors may wish to include exercises in which students reflect on when and how AI helps or harms their learning, thinking, and morale, and how it changes or expands the work they produce.

Finally, when allowing or encouraging AI use, instructors should make sure students understand that LLMs can fabricate facts and citations, code generation tools can produce incorrect or insecure code, and image generation tools can produce biased, offensive, or otherwise inappropriate outputs.

AI is only a tool, and students are responsible for all work they submit, including any inaccurate, biased, offensive, or unethical content produced with the assistance of generative AI.

3.1.9. Exercise caution with AI detectors and online exam platforms

Many tools purport to detect AI use, and some are reasonably accurate when given a piece of purely AI-generated text. However, most real-life situations are more nuanced: If a student uses AI to complete an outline or edit a section, these tools are unlikely to uncover it.

No doubt the tools will improve. **Nevertheless, we recommend against relying on AI detectors, as it risks an arms race in which students respond to automated detection by using increasingly powerful “AI humanizers” to remove signals that AI detectors are cued to catch. The result: a lot of effort on both sides that in the end serves no one.**

AI detection systems may also mistake the writing of non-native English speakers or neurodivergent students for text generated by AI. Even low rates of false positives can put students on edge and cause serious individual consequences. More generally, stepping up “policing” around AI use builds an adversarial atmosphere of distrust between instructors and students, which understandably hurts students’ motivation and morale.

Some instructors have asked if MIT will provide so-called “lockdown” browsers for conducting evaluations of students. These are online testing environments that take over the computer during an exam, preventing access to any online resources outside of those provided as a part of the test. **The committee recommends that MIT study such tools, but notes that the current generation is buggy, error-prone and feels like surveillance. At least for now, in-person proctored exams – the current norm at MIT – are a better choice in most cases, though they will require appropriate in-person spaces.**

Alternatives to these heavy-handed technological solutions are more likely to build trust and support student learning. Options include working on writing in class by hand for early ideation/drafts, creating regular project deadlines or meetings to show project progression, preventing massive amounts of work from piling up near the final deadline, and allowing instructors to give feedback at various stages (not just on the end product).

Other technical tools can also be helpful. For example, instructors can require students to do their work on platforms that capture a history of versions, and to submit the history along with their work. This can provide useful process evidence – for example if a student were to submit an assignment within a few minutes though comparable work would typically take hours. And students themselves often find these tools useful for reflecting on how their ideas have evolved.

Preparing our disciplinary system

While we hope students will follow AI policies for their own benefit, MIT needs to be better prepared to handle serious violations. MIT’s existing academic integrity framework, and the Committee on Discipline (COD) processes that enforce it, were built on a model of authorship in which the work students submitted was either their own or someone else’s. Generative AI blurs the lines. A student who uses an AI tool to brainstorm an approach, debug a function, tighten a paragraph, or generate a full draft is not “copying” in any traditional sense, and faculty in our listening sessions consistently reported uncertainty about where the line now lies, both in formulating their own course policies and in deciding which cases to refer to the COD.

MIT needs to make sure that Institute policies clarify what evidence would be required to bring an academic integrity case forward when AI is involved; the COD itself does not consider AI detector output alone sufficient.

Students should not feel policed. Durable change will require instructors to be as clear as possible about their expectations and students to understand AI misuse as an unacceptable deviation from

shared peer norms and community values rather than a violation of an arbitrary bureaucratic rule.

3.1.10. Support responsible experimentation in the curriculum

There are many unknowns about how best to integrate AI into the curriculum and how to design experiences to avoid its pitfalls. We encourage instructors and departments to pursue coordinated curricular experimentation in both these areas. It might also make sense to offer variants, for example by developing both rigorous "AI Light" or "AI Heavy" classes or pathways through the curriculum.

Some students and instructors prefer not to use AI (for a variety of reasons). We encourage instructors to take such student reluctance seriously and, where the subject matter allows it, to suggest a pathway that keeps the use of AI to a minimum.

Creating learning experiences that meet clear but perhaps evolving learning objectives, incorporating AI where applicable, and always acknowledging its presence – all while maintaining high standards – won't be accomplished by superficial tinkering. Instead, it will require bold, imaginative systemwide experimentation and evaluation. Our work on this committee demonstrates the shared will among students, instructors, and staff from every corner of the Institute to creatively address this epochal challenge together.

The committee heard from many instructors that the administrative process for making even small changes to curriculum requirements is time-consuming and burdensome, making it hard for MIT to respond to, let alone anticipate, fast-moving, far-reaching disruptions.

We urge faculty governance to review its processes this fall to ensure the presence of efficient pathways for curricular experiments that enable more rapid exploration, paired, of course, with thorough evaluation of the results. It will also be critical to raise awareness of new and existing pathways across departments, sections, and programs, so they can immediately begin exploring AI-aware substitutions and alterations in their curricula as soon as possible – or we will be left behind. At the same time, of course, it's essential to maintain the kind of long-term coherence across our curriculum that has been a hallmark of MIT education (see Section 3.3 below).

3.2. Center people, community, and the residential experience

AI is disrupting the social fabric of our campus – and the wider world. The backlash against AI in commencement speeches in the summer of 2026 was just one expression of the discontent many feel about AI itself and the role Big Tech now plays in our society.

An MIT education has always meant being challenged by a deep intellectual community that champions creativity and rigor. We must preserve the transformative power of that experience for our students.

AI has arrived at a moment when the social foundations of campus life are already under strain. Across higher education, students report¹³ high levels of anxiety, depression, loneliness, and disconnection. Confidence in institutions has weakened, political polarization has made campuses more fragile, and the post-pandemic recovery of in-person communities remains incomplete. AI is therefore not arriving in a socially neutral environment. It is imposing new pressures on a community still working to restore the foundational habits and attitudes that residential education

¹³https://healthymindsnetwork.org/wp-content/uploads/2024/09/HMS_national_report_090924.pdf

depends on: showing up, building trust, enthusiastically joining shared efforts, and living with an easy sense of mutual obligation. AI could accelerate this erosion. Or it could be the impetus for a deliberate rebuilding.

The challenge for MIT is therefore neither to preserve existing practices unchanged, nor to reject AI tools that may support learning and research. It is to act decisively to ensure that AI augments the human communities through which education happens, rather than quietly replacing them – so that the machines serve us and not the other way round.

MIT should use this moment to ask what forms of presence, collaboration, mentorship, and shared intellectual life are essential to residential education – and then design AI policies, pedagogies, and research practices that strengthen them rather than allow their erosion, replacement, or collapse.

3.2.1. Define and communicate the value of residential education

As a residential institution, MIT needs to clearly define the value of bringing students to campus and engaging avidly in face-to-face discussions, dialogue, and problem solving. No one reading this report was educated without these basic human interactions. But we cannot take for granted that incoming students will appreciate their value.

This is the core of MIT: In the spirit of “Mind and Hand,” we work together on real projects that require deep thought and careful, rigorous work.

It’s clear that unthinking, “opportunistic” use of AI tools to complete assignments diminishes what students learn and damages the learning communities they belong to. Learning works when it’s both challenging and social; knowledge is built through cognitive friction, whether that’s disentangling the steps of a mathematical proof with your study group, adjusting an experiment over and over until it works, or having a spirited argument with a peer (rather than getting “the” answer from AI). *That’s* why it matters for students to go to college!

Students need to believe that doing the work themselves is in their self-interest. It will take more than top-down AI use policies to inspire that conviction. As a community, we need to cultivate a campus culture in which (1) AI restrictions that instructors impose clearly support student learning, personal growth, and professional success; and (2) students have a shared culture that values and fosters residential learning and discourages excessive reliance on AI.

MIT’s official mission¹⁴ and values¹⁵ unite us all. But our daily decisions, conscious and unconscious, are largely shaped by the unspoken norms and values of our immediate communities – classes, study groups, departments, labs, offices, teams, clubs, living groups and more.

For our students, it’s these in-person communities that foster the social relationships – formal and informal – that allow them to grow, thrive and prepare for their lives beyond MIT.

Employers prize MIT graduates for excellent problem-solving and technical skills. But they also want people with the skills of human connection. By prompting students to turn to AI instead of to their fellow students, TAs, or professors, these technologies offer the illusion that human interaction is inconsequential. But learning with and from other people simultaneously supports learning and builds the human skills that make the world go ‘round.

¹⁴<https://web.mit.edu/about/mission-statement/>

¹⁵<https://web.mit.edu/values/>

3.2.2. Strengthen social connection and personal wellbeing

AI threatens the shared social fabric of MIT and creates new risks to the well-being of our students. MIT needs to invest in creating shared norms, establishing new community rituals, and addressing new academic, social, and emotional challenges introduced (or exacerbated) by AI.

While our core recommendations center on policies for the classroom and AI training, we encourage the Institute to take many other steps to build the kind of community that AI can't replace. (The recent report of the MIT Task Force on the Undergraduate Academic Program¹⁶ includes similar recommendations.)

In the past, the technical General Institute Requirements (GIRs) provided not only a shared academic foundation and problem-solving skill set for first-year undergraduates but also an intense bonding experience around a common core. But this core is fracturing as students rush to take major-specific classes in their first year and put off some of the GIRs until senior spring. Students' other shared rituals – campus preview weekend, orientation, convocation, the brass rat, and commencement – are widely spaced milestones. Outside of their research groups or labs, graduate students have even fewer shared rhythms and rituals.

Anecdotally, it appears that the rapid adoption of AI may already be affecting student well-being and creating new long-term challenges. Some students are becoming dependent on AI for schoolwork or emotional support; losing their sense of agency, purpose, and confidence; and worrying about their career prospects in an AI-saturated world. Instructors and other academic staff should be prepared to recognize AI-related academic, social, and emotional challenges and to direct students to appropriate support.

MIT should continue to invest in and expand its first-year learning communities programs¹⁷. But we also need to engage many more people across campus in conversation about shared norms and why community matters. One example could be a series of panels that bring together students, people from industry, instructors, and staff, to talk about where they come from and what they love: not just what they do, but how they have handled obstacles and made hard choices in life and career. Other ideas include campus-wide celebrations of human skills and accomplishments, "Tech Free Times" (periods during which MIT would not schedule classes, office hours, or meetings and instead support activities centered on in-person connection), and broadening MIT Reads¹⁸, perhaps at the department or residence level, to encourage more people to participate in collective reading, reflection and discussion.

3.2.3. Encourage instructor disclosure around their own AI use

Just as MIT needs norms around students' AI use, the community also needs to agree and adhere to norms on the appropriate use of AI by instructors. If instructors plan to present students with content substantially generated by AI, or to use AI for some aspect of evaluation, grading, or feedback, we strongly recommend that they be transparent with their students about how and why AI is being used.

While AI can be a remarkably useful tool to help instructors improve lecture notes, generate slides, and create animations or simulations, students made clear that its over-use in the classroom can

¹⁶<https://gue.mit.edu/tfuap/>

¹⁷<https://firstyear.mit.edu/first-year-advising/learning-communities/>

¹⁸<https://libraries.mit.edu/mit-reads/>

make them feel that teaching is not a priority for their instructors. A common refrain: “Why should I bother coming to class or doing the work if the teacher is just going to give an AI-generated lecture?” In a poll¹⁹, most students told the Tech they are uncomfortable with the idea of AI teaching assistants.

Our listening sessions reinforced how alert students are to instructors’ behavior: Students notice when instructors are clearly using AI for slides, feedback, or grading while imposing restrictions on student use, and they perceive it as a double standard. This points not only to a risk but also to an opportunity: by modeling responsible use of AI, instructors can help reinforce the new social contract our community will need to adapt to an AI-infused world.

With respect to evaluation, while AI can be a good way to check assignments against a rubric or automate some aspects of evaluation, many students also have a justifiably strongly negative reaction if, after investing their own time in an assignment, the only feedback they receive is from a machine. An alternative way to use AI graders is to provide them directly to the students along with the assignment, as a way for students to get useful feedback on their performance rather than a way to determine final grades. If we view assignments as a critical tool for learning, both the instructor and the student need to invest their own thoughts and efforts in the process.

3.2.4. Teach effective, responsible, and ethical use of AI

In the fall 2025 Tech Survey, more than two-thirds of students who responded felt that AI would be important in their careers, yet only 25% felt that MIT was adequately preparing them to use AI.

The effective, responsible, and ethical uses of AI are distinct – but interrelated. Students need ways to build strength in each area. We recommend immediately integrating these topics into undergraduate and graduate orientation, and infusing them throughout the curriculum as soon as possible, so community members can build a shared sense of competence in assessing and evaluating AI systems.

- To use AI effectively, students need to learn how to specify a problem or prompt, how to verify an output, when a model is likely to hallucinate, and when not to reach for AI at all.
- Using it responsibly begins with grasping the difference between augmentation and automation when it comes to education: when AI helps a human do better work, versus when it replaces the human (and thereby eliminates the learning that would come from the human doing the work). Responsible use also entails knowing how to disclose AI’s contribution honestly across coursework, theses, internal reports, and published research.
- Ethical use requires each of us and all of us together to explore issues our community has not yet worked through, including judgments about the provenance of training data; the potential for bias; the risk of homogenization in outputs; environmental and resource cost; intellectual property issues; and the line between assistance and misrepresentation of authorship.

These themes can be instilled through the curriculum in a range of ways: Capstone projects in a given major could include AI-related components, such as opportunities to reflect on the limits of AI and its appropriate use within the discipline. Departments may also choose to develop “AI-I-M” (AI Intensive in the Major) courses, and similar foundational courses for grad students, to teach students how the use of AI meshes with practices in their discipline. Communication Intensive

¹⁹<https://thetech.com/2025/11/25/llm-survey-results>

subjects should also provide instruction and practice around how to use and when to avoid AI in writing, including the risks of hallucinated citations and of the homogenization of language and ideas. Graduate advisors and research groups should incorporate discussions of appropriate use of AI in conducting and publishing research.

3.2.5. Recognize and mitigate negative impacts of AI

As we learned from both listening sessions and survey results, many people at MIT object to the way most AI models were trained – by reading vast amounts of data from the internet, often without the authors' explicit consent. Some community members decline to use AI at all because they object to how the data has been harvested without benefit or attribution to its creators.

MIT needs to recognize the broad range of attitudes and perspectives around AI, from enthusiasm to skepticism to outright rejection, and create educational experiences that will allow all students to thrive academically and personally. MIT should also encourage research and policy efforts to mitigate the negative impacts of AI.

While corporations aggressively push their AI products in the education sector, including on our own campus, people are becoming more concerned with the growing size and power of leading technology companies. This is particularly true for students, who have grown up under the sway of powerful social media firms. They would like to see both greater responsibility and greater transparency around the societal, environmental, and social impact of AI tools.

3.2.6. Acknowledge AI use in theses and other research work

Student research is a hallmark of an MIT education. Many of our students contribute to research publications, and student theses, made publicly available through the MIT Libraries' MIT Open Scholarship²⁰ repository, share our research with the world. Because these publications not only disseminate knowledge but also launch careers, they should reflect the value and integrity of an MIT education.

All theses should include a statement about how AI was used in the production of the thesis. AI should never be listed as a co-author. Most journals and conferences also require acknowledgment of AI use, and some prohibit certain use cases, so researchers should check guidelines before preparing manuscripts for publication. The MIT Libraries provide guidance on how and when to cite AI tools²¹.

Just as in the classroom, norms play an important role in how AI is used in research, particularly in the preparation of manuscripts and theses. What uses of AI in research are acceptable? How do we educate researchers in the practice of discerning when AI output is useful and factual? **Standards and practices in different communities will vary. But we believe that all MIT researchers should view AI as a tool under the direction of a human who is responsible for verifying the accuracy of its output.**

3.3. Build processes, teams and tools for continuous reflection, iteration, and improvement

The previous sections outlined many changes, both immediate and longer term, that MIT will need to make to adapt to AI. Given the startling rate of change in AI, however, the solutions we offer

²⁰<https://dspace.mit.edu/communities/6fc02cc2-0d14-4023-8a6f-d9900d0c4302/search>

²¹<https://libguides.mit.edu/cite-AI-tools>

today must continuously evolve. Our recommendations in this section are intended to address this need for deliberate ongoing adaptation.

3.3.1. Establish an ongoing AI and education committee

It was clear to all of us on the committee that, although our report is finished, MIT's work on this subject has only just begun.

To put our recommendations into practice, MIT should establish an ongoing committee. It will serve to monitor the on- and off-campus impacts of AI; evaluate changes in AI models and technology; and incorporate evolving evidence from the economy, learning sciences, and campus experience into policies and practices.

The committee will be the primary group responsible for strategy, monitoring, evaluation, and policy coordination, working closely with the AI Leads, the AI Fellows, and the AI Implementation Team (whose roles we describe in the sections below).

3.3.2. Create school/college- or department-level AI Leads

We recommend appointing AI Leads in each school and the college (or possibly each department) to help with local AI curricular planning, adaptation, and adoption.

Their responsibilities could include shaping tailored AI policies, planning how to adapt each unit's course offerings in light of AI, and developing new unit-level education and assessment models such as competency-based learning, portfolios, or both.

3.3.3. Fund AI Fellows and an AI Implementation Team

We also recommend that MIT fund an AI Implementation Team and AI Fellows – staff members, postdocs, UROPs, and graduate students with the skills to help instructors adapt their courses to AI in ways that make the most of the latest AI technology and are backed by the principles of learning science.

Including experts in technology, AI, and learning sciences, the Team and the Fellows would work at three levels: supporting individual instructors; developing campus-wide tools; and serving as implementation experts with an "ear to the ground" to study how these tools are being used across campus and to what effect.

Transitioning from longstanding classroom practices and assessments to a new set of AI-aware or AI-enabled practices and assessments presents a daunting task for any individual instructor. Doing this at scale across the Institute will be a massive undertaking.

If we expect instructors to make more than incremental improvements, we need to offer them guidance, direct support, and community. Given that existing instructor and curriculum support units like the Teaching and Learning Lab are already at capacity, the success of this endeavor will require expanding the pool of expert staff.

3.3.4. Create an AI Pilot Fund

Curricular changes are not easy, and many of our instructors, particularly in the largest classes that most immediately need to adapt to AI, already feel overworked and unsure of how to proceed. To

support AI-aware revision of course objectives and methods across campus, MIT needs to provide substantial resources.

We recommend that MIT and departments/schools/the college create an AI Pilot Fund that instructors can apply to for resources (such as AI credits, TAs, UROPs, and summer support) to do AI-enabled projects, explore the impact of AI on pedagogy, and create deliberately AI-free experiences.

3.3.5. Provide ongoing training and instructor support

In addition to material resources, instructors need ongoing opportunities to learn how AI is changing teaching, learning, and their own work. Many instructors described having relatively limited experience with AI and feeling unsure about how students are using it and how to use it themselves. Others told us that using AI has helped them develop interactive learning tools, create animations for slide decks, polish lecture notes, and reduce various aspects of course management busywork, while also creating new learning experiences for their students. The tools have already progressed to the point that even instructors with no background in developing software can produce impressive results.

To help instructors navigate both the risks and the opportunities, we recommend that MIT follow the lead of MIT Sloan²² and develop a set of resources for educators, including:

- Holding regular “lunch-and-learn” style seminars where instructors can hear from colleagues about how they are using and adapting to AI. Instructors across the Institute were eager for such “communities of practice” both within and across disciplines.
- Creating online and in-person training on the use of AI for teaching and administrative tasks, such as how to develop animations and simulations, use coding tools, build task-tracking systems, and so on. While these might incorporate components on ethics and AI safety, we urge MIT to avoid “checking the box” by adopting generic third-party AI training programs and instead invest in ways for our own community to provide these resources.
- Institute-sponsored annual or bi-annual workshops on AI and Education for both the MIT and wider academic community to convene and discuss developments and best practices.

Some of these changes can be facilitated by the AI Implementation Team and Fellows mentioned above. The Teaching and Learning Lab, which focuses solely on supporting residential education, can play a key role in helping to organize and conduct some of these activities.

3.3.6. Develop metrics

To understand how AI is reshaping our community and classrooms, **MIT should begin tracking metrics around AI use, campus engagement, student satisfaction, and post-graduation feedback.** This could include additional AI-focused questions in the biennial MIT Quality of Life survey, or perhaps other metrics, such as additional questions in course evaluations about student/staff engagement and AI use.

²²<https://mitsloanedtech.mit.edu/ai/>

3.3.7. Ensure equitable technology access

MIT should continue to provide a “model agnostic” AI platform such as Parley, that does not tie our community to one model provider, and ensure that this platform provides the level of access that our community needs.

Different generative AI models vary significantly in their performance and behavior on different tasks. Commercial AI providers like OpenAI, Google, and Anthropic all offer for-pay hosted commercial services, with the highest tier of plans costing as much as \$200/month (as of June 2026). Some researchers reportedly spend thousands more per month on pay-per-invocation requests (“API Access”) for coding tools and other use cases. These commercial providers continually renew their models to offer state-of-the-art performance.

For the MIT community, the result is an uncomfortable inequity: Many students here can spend \$200 or more per month out-of-pocket to access cutting-edge AI systems; many of their peers cannot. Some students can thus access tools that can complete substantial course projects or solve very difficult math and coding exercises overnight; others must make do with less advanced models, even if course policy allows unrestricted AI use. This imbalance has the potential to create enormous disparities in course performance.

Similar concerns apply outside the classroom, to administrators and researchers whose productivity is being measured against colleagues with access to more or less powerful models.

As an important first step in addressing access and privacy concerns, IS&T offers the Parley System²³, which gives everyone at MIT access to a range of different models, both commercial and open, and provides each user with up to \$30/month of free credits to use them. API access for use with coding tools was released in the Summer of 2026.

The access Parley provides is very helpful, but it may not be sufficient for all MIT classes, particularly in software engineering where industry is rapidly adopting these tools. Some users have also expressed concern that the \$30/month “budget” for each user is too low. We recommend that MIT continue to survey the community and consider providing expanded access for certain types of uses (e.g., specific AI-focused classes) where students need additional resources. Additionally, we have heard from several groups that even now, they still pay for commercial models because subscription plans effectively subsidize token usage relative to Parley’s per-token pricing.

Because agentic systems and tools are increasingly part of AI practice, whatever systems we employ should provide access to agentic coding and working tools, in addition to chat interfaces. We also need to ensure that we provide adequate access to computing resources that students require for classes and projects where they are training or fine-tuning models, considering options of both on-campus and cloud-hosted resources.

3.3.8. Protect sensitive data and preserve model choice

We recommend that MIT continue to invest in access to state-of-the-art tools that allow students, instructors, and staff to get work done, without having privacy concerns about the chats, data, or applications they’re using. MIT must take responsibility for ensuring this layer of privacy while remaining model agnostic.

One concern about commercial AI tools is that sensitive data may be incorporated into their training sets. Most plans, especially paid ones, allow users to deliberately disable the capture of their data

²³<https://parley.mit.edu>

for training purposes – but the default is often to record user sessions for training. This leads to the possibility of private, sensitive data from MIT users finding its way into the answers that these models provide.

MIT's IS&T provides a good reference around allowed use²⁴ for different types of data with these commercial AI providers.

Some students and faculty have concerns about how their publications (including student theses, journal articles, and books) might be used to train AI. The MIT Working Group on Scholarly Content and Generative AI²⁵ has produced guidelines and resources on this topic.

We strongly support Parley's model-agnostic approach, and we urge MIT not to commit to any single commercial AI ecosystem, while maintaining first-class access for all. We would also like to see MIT offer more hosted access to local AI models where users can issue requests without fear of training data capture by commercial providers. Recent open-source models that can be run on local GPUs provide capabilities similar to commercial models on many tasks and will continue to improve, but to run locally, these largest open models require very large GPU clusters.

3.3.9. Establish privacy, logging, and auditing policies

AI systems are not just tools for working. Since the very earliest days of the technology, people have used them for emotional and psychological support. This even occurred with one of the earliest systems, the MIT-created ELIZA chatbot (though ironically it was mainly designed to demonstrate the *shallowness* of human–AI interactions).

Today, members of the MIT community share information with AI systems that includes the most personal and sensitive aspects of their identity and experience, posing questions about their health, relationships, emotional lives, goals, aspirations, and anxieties. Our students are part of the generation most likely to use AI for this kind of personal advice²⁶.

Because Parley is an MIT-administered system, one opportunity it provides is to audit usage of AI models; in theory, IS&T can see every request from users to Parley and responses from the backend API models.

MIT therefore needs to establish a clear policy about what types of logging and auditing will be allowed. There are competing interests: On one hand, users may want to use these models to ask private or personally sensitive questions with the expectation that their data will be confidential. On the other, instructors may find it valuable to see how their students are using AI chat sessions or logs to see if learning goals are being met. MIT may need to explore anonymization technology, or establish a practice so that classroom assignments where logs are shared with instructors are clearly labeled and advertised.

A few questions MIT should resolve: How should we retain and store data? How should we monitor (or intervene in) AI interactions when users raise the possibility of harming themselves or others, or otherwise show signs of psychological distress? How can we best communicate these policies (and their implications) to the campus community? These issues need to be discussed openly and resolved transparently.

²⁴<https://ist.mit.edu/ai-tools>

²⁵https://provost.mit.edu/wp-content/uploads/2026/08/Scholarly-Content-Gen-AI-WG-at-MIT_-Recommendations.pdf

²⁶<https://www.pewresearch.org/internet/2026/06/17/americans-and-ai-2026-chatbots-smart-devices-and-views-on-impact/>

3.3.10. Monitor AI costs and environmental impact

Many community members shared serious concerns about the environmental costs of AI and the impact of data centers on communities where they're built. Issues include data centers' energy and water consumption, the impact of their construction, and the materials used to build their hardware. These have been echoed in popular literature and were the subject of a report at MIT²⁷²⁸. Both training and inference incur energy costs, and the rapid response of models, as well as the constant churn of training newer and more powerful models, has an environmental impact.

MIT should address this both through campus practice and through research.

While we believe it is essential that MIT provide access to AI tools, **we also recommend that the Institute invest in ways to provide information about the actual environmental and financial costs of using AI. This could include publicizing estimates of the environmental cost of different models on Parley, or encouraging the use of AI tools with lower impact. We also encourage MIT, to the extent possible, to audit its own AI footprint.**

There is ongoing work on campus to measure the environmental costs of data centers and find ways to reduce them. The impacts are often difficult to quantify, both because companies keep this data private and because they are diffuse, encompassing everything from construction to training costs spread over the life of the inferences. Many groups at MIT – including the MIT Energy Initiative, the MIT Climate Project, and MIT Climate and Sustainability Consortium – are exploring innovative answers around the future of data centers and communities. MIT could also contribute in important ways to designing AI models that use much less energy.

4. Conclusion

We hope this report reflects the complexity, urgency, and nuance of the challenges and opportunities that AI presents for MIT and helps the Institute community come together to take action.

AI has the potential to help our students be more productive, learn more efficiently, and build skills that can help them make a better world. What's more – though this exceeds the scope of this report – AI's extraordinary capabilities are already transforming the practice of scientific research, on our campus and elsewhere. But AI also imposes many disruptive risks to the process of learning, the practice of teaching, and the social fabric of residential education.

A strong shared view

A striking finding from our work as a committee was that – although our members came from across the Institute, including MIT's five schools and the college, and from different professional perspectives and generations – we all shared the same fundamental concerns about the risks AI inflicts on our students and our community, and the same urgent determination that MIT lead in pioneering a new model of education that celebrates human wisdom and connection.

MIT's unique position

Because of MIT's longstanding leadership in AI, its role as one of the world's preeminent institutions for research and education, and its long tradition of hands-on education, the Institute is uniquely positioned to lead higher education in this high-speed, high-stakes adaptation to the AI era.

²⁷<https://news.mit.edu/2025/explained-generative-ai-environmental-impact-0117>

²⁸<https://mit-genai.pubpub.org/pub/8ulgckc/release/2>

This is not an optional exercise. We must demonstrate how to integrate AI thoughtfully and deliberately into the classroom, the research enterprise, and the experience of residential education in ways that ensure learning, advance discovery, and prepare students for the future.

There is a lot to do.

Starting now, the MIT community must embark on a set of changes to make our educational approaches more AI-aware. That need not mean incorporating AI into everything we do, or even most things. But none of us can afford to be oblivious to the fact that AI is influencing what our students learn and how they learn it.

For students

In this work, we will serve our students by emphasizing what makes the residential experience so transformative and doing more of it; by rethinking what they need to know and revamping our assessments to match; by teaching them to use AI appropriately within their disciplines; and by helping them develop the skills, judgment, and attitudes to succeed beyond MIT in an AI-saturated world. We must foster a community that values and makes the most of the shared in-person MIT experience and that challenges students to join in the productive struggle of learning.

For instructors

The Institute also needs to make sure that instructors know how to use AI and when it's best not to. Though instructors themselves must take the lead in weighing what students need to know now and devising "AI-aware" assessments, MIT – and perhaps especially the departments – likely need to provide considerable support, coordination, and guidance. We must build "communities of practice" within and between schools, so instructors can learn from each other about how to use, adapt to, and defend against AI. This in turn requires the Institute and the departments to allow and support experimentation with, about, and around AI in the classroom, and to provide ways to capture and share those experiences.

Over time

Some changes will take time. We need to find ways to provide students with appropriate AI learning experiences from the time they arrive at MIT through graduation. We need to continually monitor and study the technical and durable human skills that industry demands, and ensure that we are preparing our students appropriately. We must deepen our scientific understanding of how AI affects learning and make sure our educational practices reflect that evolving knowledge. And while we reap the benefits of AI, perhaps especially in the realm of research, we must actively work to reduce its harms.

A roadmap to the future

With humility, we offer this report as a roadmap to help the Institute navigate this transition immediately and over time. The work will require thoughtful commitment and focused effort from the entire MIT community – students, instructors, staff and leadership. The goal: to preserve and enhance the distinctive transformative power of an MIT education.

It is not too much to say that MIT's mission depends on it. Fortunately, our work as a committee convinced us that the people of MIT are intensely interested in meeting this challenge and uniquely equipped to do it, so that the Institute can continue to develop in its students "the ability and passion to work wisely, creatively, and effectively for the betterment of humankind."

Appendices

A. Committee Process

Our committee consisted of undergraduate and graduate students, faculty from every school, and staff from relevant units, including the MIT Libraries and the Teaching and Learning Lab. Given the rapid pace of change in this area, the committee's work was brief but intense—we met weekly during the spring of '26 and conducted surveys, analysis of peer policies and practices, a number of listening sessions with the MIT community (including separate sessions for graduate and undergraduate students), as well as meetings with a range of groups working on AI on campus, including IS&T, the Committee on Discipline and the Office of Student Conduct and Community Standards, MIT Open Learning, the school councils, and a number of individuals we felt had particular insight and expertise in AI and education.

AI Use: None of the text of this report was generated with AI. We did supply early drafts of the report to ChatGPT to identify redundant / overlapping sections. Some of the graphs and statistical summaries in Appendix C were generated with Codex (OpenAI's AI coding tool).

Members of the committee also used AI to synthesize and present data and reports (e.g., summaries of other schools' published positions on AI) to inform committee discussions.

B. Sample AI Policy for Syllabi

We recommend that the campus agree on a menu of standard policies for AI use in classes, and that every class have a clearly stated policy selected from this menu (and adapted as necessary).

As described in Section 3.1.8 above, once instructors choose an option from the menu, they need to convey it clearly to students through the syllabus and perhaps on the course website, along with an explanation of the rationale behind the choice: Given the goals for what students will learn in the course, how does the selected AI policy support and encourage that learning?

The idea is to make sure students are clear about what is and is not allowed in each course or assignment and, more broadly, to help them understand what the costs and benefits are to their own growth.

Below is a proposed 4 option menu which may be applied to a whole course or specified for individual assignments. Using the traffic light icons can help make this system easy for students to decipher.

B.1. Unrestricted GenAI Use

Policy: Students may use any GenAI system for any purpose on assignments in this course.

Best suited for:

Courses where the primary assessment is done via in-class (AI-free) evaluations or assignments where GenAI use does not undermine the learning objectives, such as ambitious software-system projects that are beyond the capabilities of current tools or where evaluating and integrating AI output is itself part of the work.

Instructors may wish to ask students to include short per-assignment disclosures describing how they used GenAI, possibly including transcripts or session records when appropriate. Instructors

may also wish to ask students to reflect on how AI use did or did not help with the assignment or their understanding.

B.2. Limited GenAI Use: Support Tool Only

Policy: Students may use GenAI to support learning, brainstorming, editing, debugging, or generating explanations, but may not use it to produce full or substantial assignment solutions. This is similar to the way that a student might ask a classmate or TA for assistance. Assignments in this category should clearly specify which stages, processes or components allow and disallow use.

Prohibited uses may include:

Copying full problem prompts into GenAI tools, asking GenAI to solve assigned problems, or submitting AI-generated text, code, equations, or analysis as one's own.

Best suited for:

Courses where independent problem solving is central, but where AI can reasonably serve as a tutor, editor, or study aid. Students must clearly distinguish their own work from AI-assisted work and cite AI-generated material or ideas. This can be good for courses that already have clear collaboration norms and want a simple, intuitive AI policy aligned with existing academic-integrity expectations.

Note that in our experience students and faculty may differ as to what “support” means — for example, some people feel that using AI to write an outline or AI to generate diagrams is more than “support.” For instructors with strong opinions about this, policies should clearly state what types of uses are and are not allowed.

We expect that this is the appropriate policy for many of our courses, but also requires the most thought in specifying boundaries.

B.3. Required GenAI Use

Policy: Students are required to use GenAI to complete the assignment in the manner specified.

Students must use the designated tool or class of tools, follow the prescribed workflow, and submit any required documentation of AI interactions.

This may accompany one of the other categories, as there may be components that require AI and others where it is disallowed.

Best suited for:

Assignments where learning objectives include AI-assisted programming, prompt engineering, critique of model outputs, human-AI collaboration, or evaluation of AI systems.

B.4. GenAI Use Strictly Prohibited

Policy: GenAI may not be used in any form for the course.

Because policing out of class use as a reference aid or support tool is essentially impossible, this policy is difficult to enforce reliably and may create risks of both undetected violations and false accusations.


C. Survey Results

In this section, we summarize the results of three recent surveys of the MIT community around AI: a survey our committee designed in spring 2026, a set of questions we added to the MIT Quality of Life Survey (also conducted in spring 2026), and a survey conducted by the MIT Tech (mostly focused on undergraduate and graduate students) in fall 2025.


C.1. Spring 2026 AI Usage and Attitudes Survey

We conducted a survey of the MIT faculty, instructional staff and students regarding their AI usage and attitudes. We received 1632 responses, representing a 12% response rate. The response rate was fairly uniform across schools, varying between 20% (SHASS) to 10% (Sloan). The student response rate was 12%, with instructors and faculty responding at a 16% rate. We summarize a few of the key results.


Overall Usage (Never/Occasionally/Often/Very Often). The figure below shows that ChatGPT continues to be the most widely used system, with 44% of respondents (46% of students and 35% of instructors) saying they used it “Often/Very Often” (green bars.) Overall, Sloan reported a much higher rate of usage of all tools than any other group.




Faculty/Instructor Usage (No [Orange]/Not Applicable/Yes [Blue]). We asked instructors how they used AI; many use it for administrative tasks and helping prepare assignments or lecture notes, with much lower numbers using it for grading or individualized feedback.




Student Usage. Students use AI for many purposes, with large fractions reporting using it for coding, information search, writing revision and summarization. Fewer use it for whole-cloth generation of new text. As with overall usage, Sloan students use AI for all purposes at a significantly higher rate than other groups.




AI Attitudes (Strongly Disagree [Brown] / Somewhat Disagree / Neutral / Somewhat Agree / Strongly Agree [Dark Green]). Respondents broadly feel that they need AI skills and that AI makes them more productive, but also that they are becoming more dependent on it. Responses are more mixed on whether it helps learning or whether it is desirable in class. Most people do not like to use it as a writing tool.



Research Use (Never Acceptable [Brown] / Sometimes Acceptable / Unsure / Always Acceptable [Purple]). We asked what types of AI use the community felt were appropriate. Most felt that spelling and grammar, research, and code generation were acceptable uses, while most felt that using it for peer review or drafting paragraphs from outlines or notes was unacceptable. Other uses had mixed views, reflecting the difficulty in creating a uniform policy for research usage.



AI Concerns. (Not a Consideration [Purple] / Minor Consideration / Major Consideration [Dark Green]). Finally, we asked about concerns with AI use. Many users are concerned with hallucinations, but ethics and environmental impact are also of concern for many users.



Approximately 10% of respondents (N=163) wrote in a concern for the "Other" response here. The most common concern (N=79) was the potential impact of AI on learning and cognitive development. Respondents expressed concern that reliance on AI could reduce critical thinking, writing ability, creativity, reasoning, and long-term skill development, or otherwise diminish the educational value of coursework. A second broad category (N=34) reflected more general philosophical or societal concerns about the role of AI in education, including worries about the replacement of human expertise, the purpose of higher education, and the broader societal implications of AI. Other concerns included the usefulness, reliability, or educational value of current AI systems (N=20); the practices and business models of AI companies (N=14); privacy, security, copyright, and intellectual property (N=10); academic integrity and authorship (N=7); the cost of AI tools (N=5); and potential economic or labor-market impacts (N=4). Overall, although the concerns were diverse, they were dominated by questions about how AI may affect student learning, cognitive development, and the educational mission of the Institute.


C.2. Spring 2026 Quality of Life Survey

The Spring 2026 campus-wide Quality of Life Survey²⁹ also asked questions about generative AI usage and attitudes. We also briefly summarize these results here. This survey was given to all members of the MIT community, with about 8200 respondents from the main MIT campus (Lincoln Labs was also surveyed but their data is omitted in the analysis below). For this analysis we group faculty and instructors together, postdocs and grad students together, and service and support staff together.

Frequency of AI Tool Use. We asked respondents how frequently they used generative AI tools. Overall about 40% of respondents said they used these tools often or very often, with 46% of undergrads and 60% of undergrads and graduate students reporting at least often use. Service and support staff had a markedly low level of use, with 53% reporting they never used AI tools.

How often do you currently use generative AI tools (e.g., ChatGPT, Copilot, Claude, Gemini) in your work or coursework at MIT?

Main Campus only. Response range: Never -> Very often.




Generative AI Attitudes: We asked respondents a series of questions about AI attitudes: whether AI made them more pessimistic or optimistic, whether it made them feel more replaceable or capable, whether it made their work better or worse, whether it made them more or less efficient, and whether AI was unreliable or reliable.

²⁹<https://qol.mit.edu/>

Pessimistic vs Optimistic. Overall, only 23% of respondents are optimistic about generative AI (rating a 4 or 5 on a pessimism to optimism scale). Instructors and grad students/postdocs report a slightly higher level of optimism.

How do you feel, generally, about generative AI along these dimensions: Makes me pessimistic about the future (1): Makes me optimistic about the future (5)


Main Campus only. Response range: 1 -> 5.



Replaceable / Capable. The community is largely divided on the replaceable/capable axis. Overall, 40% of respondents said they felt AI made them more capable, and 27% said they felt it made them more replaceable. Postdocs, grad students, instructors, and admin staff reported that AI made them feel more capable at a significantly higher rate than undergrads and service/support staff. Notably, undergrads reported feeling replaceable at a higher rate than capable (40% vs. 34%).

How do you feel, generally, about generative AI along these dimensions: Makes me feel replaceable (1): Makes me feel more capable (5)


Main Campus only. Response range: 1 -> 5.



Worsens / Improves Work. Respondents reported that they felt AI improved their work more than they felt it worsened their work (42% selecting a 4 or 5 on the work worse / work better axis vs 21% answering a 1 or 2). Differences between subgroups were not large, although undergrads reported a slightly higher rate of “worsens work”.

How do you feel, generally, about generative AI along these dimensions: Makes the quality of my work worse (1): Makes the quality of my work better (5)


Main Campus only. Response range: 1 -> 5.



Inefficient / Efficient: There was strong agreement that people felt AI made them more efficient, with 60% of respondents selecting a 4 or 5 on this axis. The service / support group reported a lower rate of “makes my work more efficient”; otherwise differences between groups were small.

How do you feel, generally, about generative AI along these dimensions: Makes my work less efficient (1): Makes my work more efficient (5)


Main Campus only. Response range: 1 -> 5.



Unreliable / Reliable: Significantly more respondents felt that AI was unpredictable and unreliable than felt it was trustworthy and reliable, with 39% selecting a 1 or 2 vs 21% selecting a 4 or 5 on this axis. Again the differences between subgroups were small.

How do you feel, generally, about generative AI along these dimensions: Current systems are unpredictable and unreliable (1): Current systems are trustworthy and reliable (5)


Main Campus only. Response range: 1 -> 5.



MIT Guidance. We also asked respondents whether they thought MIT had given them adequate guidance about AI use. Undergrads largely reported that they had received such guidance, with 72% reporting they somewhat or strongly agreed, with other groups reporting much lower levels of adequate support.

MIT has provided me adequate guidance on how and when to use generative AI in my work or coursework at MIT.

Main Campus only. Response range: Strongly disagree -> Strongly agree.



C.3. Fall 2025 Tech Survey

In the Fall of 2025, the Tech (MIT's student-led newspaper) also conducted a survey of student attitudes and usage of AI. The results are described in detail in a recent tech article³⁰, which we summarize briefly here. The survey received responses from 1,002 affiliates, including 659 under-

³⁰<https://thetech.com/2025/11/25/llm-survey-results>

graduates, 248 graduate students and postdocs, 18 faculty, and 72 staff. Among undergraduates, 46% reported using LLMs daily and 30% several times per week; among graduate students and postdocs, the corresponding figures were 46% and 35%. The survey found that the most common use among undergraduates was explaining course material (over 80%), followed by coding assistance (around 70%), coursework completion (over 50%), and summarizing papers (over 50%). Among graduate students and postdocs, the most common use was programming assistance (88%), followed by brainstorming ideas (55%), summarizing papers (52%), and writing essays (48%). The survey also found that 35% of undergraduates reported that LLMs save them “a lot” of time and more than 45% reported that they save “a bit” of time. The survey found that 90% of undergraduates were somewhat or very concerned about overreliance on LLMs, including 67% who were very concerned. In addition, 45% were very concerned about inaccurate or misleading outputs.

The survey found that 75% of undergraduates believed faculty expectations regarding AI use were clear, while 50% believed decisions about AI use in coursework should be left to individual instructors and departments. Although 70% agreed that AI proficiency will be important in their careers, only 25% believed MIT is preparing students to use AI professionally.

D. Additional Resources

See <https://aihub.mit.edu> for more resources about AI in education and research and how we are beginning to implement the findings of this report.