Report of the Working Group on Artificial Intelligence in Administration and Operations

May 16, 2025

Report of the Working Group on
Artificial Intelligence in Administration and Operations
Table of Contents

Executive Summary

The Working Group on Artificial Intelligence (AI) in Administration and Operations was convened in September 2024 at the request of Provost Cindy Barnhart and Executive Vice President and Treasurer (EVPT) Glen Shor. The group was charged with understanding the fast-evolving landscape of AI and generative AI tools and how these technologies can be applied ethically and effectively to enhance MIT's administrative and operational processes.

Key objectives included:

- Developing guardrails and comprehensive guidance to support MIT administrators in identifying pitfalls and effectively adopting new AI technologies into their workflows.
- Benchmarking AI adoption in business operations at peer institutions
- Cataloging existing AI-based tools used at MIT across administrative, academic, and research areas
- Developing core principles and policies to guide AI use, focusing on security, privacy, and ethical considerations
- Supporting experiments and pilots to explore AI applications, using discovery sprints as a methodology to deepen insights *(led by AI Working Group 2: High Impact Opportunities)*

Between September 2024 and April 2025, the group held regular meetings and engaged subject matter experts to gain their understanding of AI's potential and risks. Highlights of expert presentations include:

- Analysis of Large Language Models (LLM)/Generative AI pitfalls and risks
- Technical implementation insights from MIT's Admin ChatBot (currently in beta)
- Strategic perspectives from peer institutions
- Research on AI's impact on workplace productivity
- Insights on AI integration in communication
- Existing use cases of AI use at MIT

Drawing on insights from the meetings with subject matter experts within and beyond MIT, this report provides detailed findings and recommendations that are organized into two parts:

- **Part 1: Foundational AI Guidelines -** addresses the opportunities and challenges of general AI use across MIT
- **Part 2: Specialized AI Solutions-** provides additional recommendations and insights regarding purpose-built AI solutions, both developed internally at MIT or provided from third parties

This report also includes a decision guide that helps MIT faculty and staff in evaluating when to leverage AI tools, along with core principles to guide AI usage across the Institute, and guidelines for labeling AI-generated content.

Key Recommendations

The rapid adoption of AI tools across MIT has created both opportunities and challenges that require thoughtful consideration. Below, the Working Group captures key recommendations to guide MIT's approach to AI implementation and oversight. While the complete list of recommendations is captured in the detailed Findings and Recommendations section, these highlight the most high-priority items.

1. While AI can help increase efficiency and automation for many administrative functions at MIT, its use must always be consistent with MIT values
2. It is important to establish a culture of accountability for AI deployments; even when AI is used to automate a function, an MIT person must always bear ultimate responsibility for the actions of the AI system
3. Maintain and continuously update MIT's existing AI guidance web page (https://ist.mit.edu/ai-guidance) as a central resource for AI tool selection, privacy considerations, enterprise options, and evolving policies
4. Provide enterprise-wide access to multiple AI tools to prevent use of unsecure, open platforms and reduce security risks
5. Ensure flexibility in AI tool selection to adapt to evolving technology and avoid dependency on single vendors
6. Provide training programs focused on:
    - Opportunities for using AI tools to increase efficiency and effectiveness for specific tasks and use cases
    - Assessing whether AI is appropriate for a given task, considering both the risks and potential benefits
    - Safe handling of sensitive and private information
    - Best practices for responsible AI use
7. Reevaluate the use of AI in situations where the risks outweigh the benefits and human oversight may not sufficiently mitigate issues (e.g. plagiarism, high-stakes decision making)
8. Promote pilot programs to test AI implementations in controlled environments, allowing teams to identify challenges and learn from failures before wider deployment
9. Establish clear guidelines for the labeling of AI-generated content; an initial set of proposed guidelines is provided on page 10. These should be updated as technology and social norms evolve, but always maintaining the principle of not misleading users about the authorship of content emanating from MIT
10. Create an additional working group to address student concerns and the impact of AI on student life
11. Be cautious about overinvesting in AI solutions at the expense of necessary improvements to other critical software needs
12. Encourage cross-sharing of emerging best practices in choice of tools and methods at MIT, highlighting what is being done, what works, what does not work, and how implementations were approached

Principles for AI Use at MIT

These principles ensure that AI use aligns with the Institute's values, prioritizing effectiveness, transparency, and human oversight while safeguarding data integrity and ethical responsibility.

1. Tool Appropriateness
    - Tool selection should prioritize effectiveness, recognizing that traditional software may be the better choice for many applications
2. Human Oversight and Individual Responsibility
    - For any use of AI in administration, an MIT-employed *responsible party* must be accountable for the system's outputs ensuring it is fit for its intended purpose and does not cause harm. Even though AI may be used to automate many internal functions, human judgement will remain essential in ensuring that AI tools are used only when appropriate and that proper precautions are taken to mitigate potential risks.
3. Content Integrity
    - AI-generated content should be reconsidered if:
        - Inaccuracies could cause harm or damage credibility, especially in public-facing contexts
        - It could undermine trust through bias or misinformation
        - It poses legal or ethical risks to MIT
4. Transparency
    - When AI plays a significant role in generating content or making decisions, it should be clearly disclosed so users understand they are interacting with AI-generating outputs. If there is no human involvement, but users might reasonably assume otherwise, this should also be stated to prevent misunderstandings
5. Data Security
    - Sensitive data, including personal, financial, pre-public, third-party and confidential information, should only be used through MIT-approved enterprise AI systems suitable for the required level of confidentiality (risk classifications can be found on the MIT Information Protection website).
    - IS&T must ensure that the list of MIT-approved AI systems is explicit and up to date about any restrictions to the use of sensitive data on those systems.

Detailed Findings and Recommendations

AI plays a growing role in administrative functions at MIT. While its potential is vast, promoting AI use involves navigating a set of technical, ethical, and operational considerations. To address these effectively, this section organizes our findings and recommendations into two parts: foundational AI guidelines and considerations for specialized AI solutions, with each section providing key findings with recommendations.

Part 1: Foundational AI Guidelines

This section addresses the opportunities and challenges of general AI use across MIT.

A. Thoughtful Tool Selection and Access

MIT faculty and staff now have growing access to AI tools that could change how they work. While these technologies can transform many aspects of their day-to-day tasks, success depends on choosing the right tools and helping our community use them safely and effectively.

Key Findings

- AI can effectively automate repetitive tasks, support decision-making, draft messages/documents, and help brainstorm new ideas
- Without helpful guidance, users may default to readily available tools without evaluating their appropriateness or security implications
- The use of non-enterprise AI tools presents significant privacy and security risks (e.g. exposing personal, confidential, financial or IP data)
- Employees would benefit from access to multiple enterprise versions of general-purpose AI tools to ensure exposure to diverse capabilities and avoid vendor lock-in

Recommendations

- Provide employees with access to multiple enterprise versions of general-purpose AI tools
- IS&T should continue to expand its capacity to provide guidance to individual units on whether AI solutions are appropriate to specific problems and on choosing versions of AI tools that comply with privacy and confidentiality requirements.
- Individual units should be given flexibility to adopt new tools as the technology evolves as long as they are consistent with the principles outlined earlier

B. Human Oversight and Accountability

The use of AI for Administration and Operations should follow the principle that while it is possible and, in some cases, desirable to delegate work to an AI system, responsibility cannot be delegated and must always reside on an MIT employee. For some use cases, (e.g. using an AI system to help compose an e-mail) it is sufficient for users to check the output of the AI system for errors or inaccuracies that could have been introduced by it, but users should be

mindful of some of the failure modes that can be difficult to catch through manual checks. For example, using AI to author public facing web pages or documents introduces a risk of plagiarism that is difficult to mitigate by manually checking the output. Uses of AI where validating individual outputs is not feasible will have to rely on extensive testing and validation to ensure that the AI system is fit for purpose and that errors will not lead to problematic outcomes.

Key Findings

- AI systems can generate convincing but inaccurate content; they have significant failure modes that makes them unsuitable for tasks that require precision and cannot tolerate errors
- Models have specific limitations in multilingual contexts due to English-centric training data and language processing
- AI can produce plagiarized content without the awareness of the person using it
- AI models can perpetuate and amplify societal biases present in their training data

Recommendations

- In the context of Administration, AI's role is to augment and support human decision-making, but final responsibility always lies with an MIT employee
- Users must take responsibility for deciding whether to use AI tools and the AI-generated outputs
- AI-generated content should be used where the risks are minimal or can be mitigated, for example through human oversight. Use cases where AI operates autonomously without human review (for example an external-facing chatbot) must be rigorously evaluated and piloted before broad deployment (See AI Decision guide)
- Encourage disclaimers for AI-generated content to promote transparency and maintain trust where necessary - such as when AI operates autonomously without human oversight. Disclosures should clearly inform users so they can make an informed choice about whether to use the output (See Guidelines for Labeling AI Content)
- Limit multilingual support to languages where there is enough expertise to directly validate the output in those languages. Specifically, it should never be assumed that just because the system's performance or output has been validated in one language the system can be trusted to work reliably in an untested language

C. Training and Support

MIT faculty and staff, as well as students, would benefit from clear guidance and resources to use AI tools effectively. Having the right training and support helps everyone make informed decisions about when and how to use these technologies.

Key Findings

- Convenient, accessible, and quality training is essential for responsible AI adoption, especially regarding topics like data privacy and other security risks

- Users might be unable to identify potential bias and ethical concerns and would benefit from clear guidance on which AI tools are appropriate for different tasks

Recommendations

- Provide training for MIT staff on the appropriate use of enterprise AI tools compared to free-form AI tools. Emphasize the benefits of MIT-provided enterprise tools, such as enhanced data privacy and reduced risks of sensitive information leakage, to ensure informed and responsible tool selection
- Discourage the use of AI tools in high-risk scenarios where errors could have serious consequences, such as critical decisions, financial implications, safety concerns, or plagiarism risks
- Continue to develop and evolve a web page that:
    - Guides faculty and staff in selecting appropriate AI tools for specific tasks
    - Explains MIT's enterprise licenses and preferred tools
    - Includes understanding data and privacy risks for AI use
    - Updates regularly to reflect new tool availability and policy changes
    - Serves as the primary source for MIT's AI tool guidance
- Form a dedicated working group focused on student AI use, support, and implementation

Part 2: Specialized AI Solutions

This section provides additional recommendations and insights regarding purpose-built AI solutions, both developed internally at MIT or provided from third parties.

A. Development and Implementation

As MIT teams create and adopt specialized AI tools to support administrative functions, we must carefully plan how to build them, maintain them, and manage their costs. Getting these systems right requires attention to testing, data quality, and budget considerations.

Key Findings

- Usage-based pricing models can lead to unpredictable costs, especially for widely adopted tools; AI API calls and credits can accumulate significant expenses if not properly monitored and managed
- When an AI tool is built by or represents MIT, it should undergo greater scrutiny to ensure it maintains trust and is free from legal or ethical risks
- Vendor solutions should align with MIT's data and security requirements
- Frequent updates to an AI tool or service can disrupt workflows
- Sensitive and private data should not be shared with third-party vendors

Recommendations

- Establish clear criteria and ongoing cadence for evaluating and testing AI outputs

- Tailor review processes based on the system's impact, applying stricter oversight to AI tools that are public facing, could undermine trust, or pose legal or ethical risks
- Factor in potential scaling costs when evaluating specialized AI tools, particularly for applications likely to see widespread adoption
- Regularly review and optimize AI credit usage to prevent unexpected budget impacts (e.g. consider implementing usage quotas)
- Create clear update and retirement procedures for data
- Ensure sensitive and private data is not shared with third-party vendors
- Assess whether enterprise implementations can isolate MIT's data to minimize exposure of sensitive information to third-party vendors
- Implement status monitoring and alerts for integrations, especially for unexpected downtime

B. AI Adoption

MIT faculty and staff should be encouraged to try AI tools through pilot programs and share what they learn with others. These experiences - both successes and challenges - help our community make better decisions about using AI effectively.

Key Findings

- Pilot programs and controlled experimentation provide opportunities to learn from failures and refine AI implementations
- Pilots should focus not just on functionality but also on cost estimation and planning; ongoing costs for supporting AI systems can be substantial
- Sharing successful use cases across DLCIs can foster innovation and enhance understanding of AI's capabilities and limitations

Recommendations

- Support and promote pilot programs to test AI implementations in controlled environments before broader deployment
- Create a small team that can develop and run training programs for AI integration into staff workflows
- Encourage sharing AI success and failures to promote a culture of continuous improvement

AI Decision Guide

This guide serves to assist MIT faculty and staff in evaluating if AI is appropriate based on its impact on people and the nature of its outputs. AI is most useful in applications that require unstructured input such as natural language, code, or images. However, for tasks that require precision and repeatable results (e.g. a calculation), AI may not be the most effective solution or cost effective.

1. **Do not use AI**
   AI should be avoided in these situations, as a non-AI solution is more appropriate:
   a. The task does not need the general capabilities that AI provides as a non-AI-based solution is more cost effective and more reliable
   b. The task requires full accuracy, and even a small probability of error would pose unacceptable risk to the institute.
   c. AI-generated content could create reputational harm, even if the output is accurate (e.g. if a heartfelt message from leadership was found to be AI generated, it would lead to reputational harm even if the message itself was unobjectionable)
   d. The task involves sensitive or confidential information, and there are no MIT-approved AI tools that ensure data security
2. **Fine to use AI, but a person should vet all outputs**
   AI may be used, but outputs should be reviewed by a person before use:
   a. Tasks that benefit from the use of AI such as brainstorming, summarization, and draft writing, but require human oversight to check for biases, errors, or other inaccuracies
   b. Outputs that help inform decision-making, but a responsible party must make the final call

   AI tools in this category must still align with MIT's policies regarding security, data, and privacy. Additionally, any non-MIT approved AI tool, including public off-the-shelf options that collect data, should be used only when the information is neither sensitive nor valuable to third parties.
3. **Fine to use AI, including without human supervision, but the system should be validated for its purpose**
   AI can be used without direct human oversight if:
   a. The risk of errors has been shown to be low through rigorous testing
   b. A responsible party has thoroughly considered the potential consequences of errors and put in place measures to mitigate their impact
   c. Users understand that content is fully AI-generated and can decide whether to use it or seek an alternative, such as connecting directly with a human
   d. The AI-generated content serves a functional purpose and is not expected to convey human intent or emotion

   AI tools in this category must still align with MIT's policies regarding security, data, and privacy. Additionally, any non-MIT approved AI tool, including public off-the-shelf options that collect data, should be used only when the information is neither sensitive nor valuable to third parties.

Guidelines for Labeling AI Content

The purpose of labeling AI-generated content is to promote transparency, maintain trust, and prevent unintentional plagiarism. In communications and published materials, the level of labeling should reflect the extent of AI's role and in the context in which the content is shared. While labeling AI generated content is important in fostering trust, labeling content as AI generated does not absolve its creators from responsibility for ensuring its quality and suitability for use.

1. **Must Label AI-Generated Content: AI-generated content must be labeled when:**
    a. It is delivered directly to a user without human review or modification (e.g. AI-generated responses in chatbots, AI-written information automatically sent to users)
    b. The accuracy or authenticity of the content might be questioned if its AI origin is not disclosed

    Example label: "*This content was generated by AI without human review.*"
2. **May Benefit from a Label: Labeling is recommended, but not required, in cases where:**
    a. AI played a substantial role in creating the output (e.g. internal decision-making, knowledge-base articles)
    b. The content is intended for public-facing audiences and transparency would enhance credibility (e.g. web content or widely circulated documents)
    c. When there is an expectation that the content was fully created by a human.

    Example label: "*This content was generated with AI assistance and reviewed by a human.*"
3. **When AI-Labeling Might not be Necessary: AI-generated content does not require labeling if:**
    a. AI only provided minor enhancements such as grammar correction, format suggestions, or stylistic improvements
    b. An AI system played a role in the process of producing a document, but not in actually producing the document itself (for example, it was used to do research for an article, but not in writing the article itself)

Benchmarking AI Maturity at Peer Schools

The Working Group conducted a benchmarking study of AI maturity across leading American R1 Institutes/Schools to evaluate the current landscape of AI use in higher education. The assessment, purposely anonymized, examined six schools across five key dimensions of AI readiness: strategic planning, governance, enterprise tools, training programs, and pilots and use cases. Our assessment revealed that while MIT is not currently at the forefront in providing comprehensive AI guidance to faculty and staff, the recommendations and next steps outlined in this report establish a clear path forward. At this time, MIT's AI maturity is developing across all key dimensions, with opportunities for growth in each category.

For a detailed analysis of the benchmarking findings, please see Appendix E: Benchmarking Insights.

AI Maturity Matrix at Peer Schools

Institution	AI Strategy & Web Presence	AI Governance & Ethics	Enterprise AI Tools	Training	Pilots & Use Cases
Peer Institute/School A	Leading	Developing	Leading	Leading	N/A
Peer Institute/School B	Leading	Developing	Developing	Developing	Leading
Peer Institute/School C	Leading	Leading	Leading	Developing	Leading
Peer Institute/School D	Leading	Leading	Leading	Developing	Developing
Peer Institute/School E	Leading	Leading	N/A	Leading	N/A
Peer Institute/School F	Developing	Leading	Leading	N/A	N/A

Maturity Legend:

Leading: Mature AI program with robust web presence and committee recommendations	**Developing**: Basic information available; work in progress	**N/A**: Unable to verify information via web search; more investigation is needed

Next Steps

The Working Group proposes the following next steps to guide MIT's adoption of AI in administration and operations over the next few years:

Enterprise AI Tools

- Continue to purchase, deploy, and promote enterprise-wide AI tools over free-form AI tools
- Encourage staff to use provided AI tools to ensure data, privacy, and security standards are maintained

Training and Support

- Launch comprehensive training programs focused on:
    - Appropriate tool selection
    - Human accountability
    - Reevaluating the use of AI in situations where the risks outweigh the benefits and human oversight may not sufficiently mitigate issues (e.g. plagiarism, high-stakes decision making)
    - Privacy and security considerations
    - Opportunities for using AI tools to increase efficiency and effectiveness for specific tasks and use cases
    - Best practices for responsible AI use
- Publish and maintain:
    - AI Principles
    - AI Decision Guide
    - AI Guidelines for Labeling AI

Policy Development

- Partner with OGC and Audit to review existing AI policies and have a process in place to ensure policies remain up to date as technology evolves
- Update vendor agreements and contracts to ensure compliance with MIT standards and ensure that the guidance provided in the list of MIT-approved enterprise AI systems is consistent with these agreements.

Additional

- Promote and support AI pilot programs across DLCIs
- Establish mechanisms for sharing successes and lessons learned
- Form a dedicated working group to address student impacts and concern

Conclusion

AI offers transformative potential for MIT's administration, enabling the Institute to operate more efficiently and effectively. However, realizing this potential requires careful planning, ethical considerations, and robust training. By adopting the recommendations outlined in this report, MIT can harness AI's benefits while mitigating its risks.

Appendix A: Table of Recommendations

Foundational AI Guidelines
A. Thoughtful Tool Selection and Access
Provide employees with access to multiple enterprise versions of general-purpose AI tools
IS&T should continue to expand its capacity to provide guidance to individual units on whether AI solutions are appropriate to specific problems and on choosing versions of AI tools that comply with privacy and confidentiality requirements
Individual units should be given flexibility to adopt new tools as the technology evolves as long as they are consistent with the principles outlined earlier
B. Human Oversight and Accountability
In the context of Administration, AI's role is to augment and support human decision-making, but final responsibility always lies with an MIT employee
Faculty and staff must take responsibility for deciding whether to use AI tools and the AI-generated outputs
AI-generated content should be used where the risks are minimal or can be mitigated, for example through human oversight. Use cases where AI operates autonomously without human review (for example an external-facing chatbot) must be rigorously evaluated and piloted before broad deployment (See AI Decision guide).
Encourage disclaimers for AI-generated content to promote transparency and maintain trust where necessary - such as when AI operates autonomously without human oversight. Disclosures should clearly inform users so they can make an informed choice about whether to use the output (See Guidelines for Labeling AI Content)
Limit multilingual support to languages where there is enough expertise to directly validate the output in those languages. Specifically, it should never be assumed that just because the system's performance or output has been validated in one language the system can be trusted to work reliably in an untested language.
C. Training and Support
Provide training for MIT staff on the appropriate use of enterprise AI tools compared to free-form AI tools. Emphasize the benefits of MIT-provided enterprise tools, such as enhanced data privacy and reduced risks of sensitive information leakage, to ensure informed and responsible tool selection

Discourage the use of AI tools in high-risk scenarios where errors could have serious consequences, such as critical decisions, financial implications, safety concerns, or plagiarism risks
Continue to develop and evolve web page that: • Guides faculty and staff in selecting appropriate AI tools for specific tasks • Explains MIT's enterprise licenses and preferred tools • Includes understanding data and privacy risks for AI use • Updates regularly to reflect new tool availability and policy changes • Serves as the primary source for MIT's AI tool guidance
Form a dedicated working group focused on student AI use, support, and implementation
Specialized AI-Based Tools
A. Development & Implementation
Establish clear criteria and ongoing cadence for evaluating and testing AI outputs
Tailor review processes based on the system's impact, applying stricter oversight to AI tools that are public facing, could undermine trust, or pose legal or ethical risks
Factor in potential scaling costs when evaluating specialized AI tools, particularly for applications likely to see widespread adoption
Regularly review and optimize AI credit usage to prevent unexpected budget impacts (e.g. consider implementing usage quotas)
Create clear update and retirement procedures for data
Ensure sensitive and private data is not shared with third-party vendors
Assess whether enterprise implementations can isolate MIT's data to minimize exposure of sensitive information to third-party vendors
Implement status monitoring and alerts for integrations, especially for unexpected downtime
B. AI Adoption
Support and promote pilot programs to test AI implementations in controlled environments before broader deployment
Create a small team that can develop and run training programs for AI integration into staff workflows
Encourage sharing AI success and failures to promote a culture of continuous improvement

Appendix B: Working Group Charge and Membership

To: Members of the Working Group on Artificial Intelligence (AI) in Administration and Operations

From: Glen Shor, Executive Vice President and Treasurer; Cindy Barnhart, Provost

Date: September 13, 2024

Subject: Working Group Charge and Membership

Thank you for your willingness to serve on the Working Group on Artificial Intelligence (AI) in Administration and Operations. We appreciate your devoting time and energy to be part of this group that will work to ensure that MIT uses AI effectively and ethically to improve our operations. This work is of utmost importance as we begin a seven-year roadmap to modernize our enterprise systems and related business processes.

The Working Group will be comprised of two workstreams that will move forward in parallel and in close coordination to expedite this work.

Setting the Foundation

The first workstream will endeavor to set a sound foundation for the use of AI in administration and operations. This includes understanding the technical complexities involved in pursuing high-impact opportunities, examining legal constraints and evaluating potential pitfalls. The group will develop guidelines for procuring and deploying AI-based systems for administration and operations and suggest policies around security and privacy. The group will work to understand current efforts in industry and at other universities that are actively experimenting with and deploying AI in their operations.

High-Impact Opportunities

The second workstream will work to determine how deploying AI-based systems could lead to high-impact improvements in administration and operations, and they will prioritize these opportunities based on need, risk and feasibility. Members will develop business cases to recommend a set of AI-based system experiments and pilots for Institute administration and operations. They will identify key areas where AI-based systems are already being used at MIT and explore how AI is used in operations outside of MIT to inform how the use of AI can be appropriately and effectively expanded at MIT. Some ways in which we already use AI in operations might be adapted by additional units/functions. Based on guidance from the first workstream, the group will ensure that all technical complexities and legal constraints are fully understood.

Given the demands on your time, we appreciate that you have agreed to serve on the Working Group and to share your invaluable expertise. We are particularly grateful to Distinguished

Professor of Computing and Associate Director of CSAIL Armando Solar-Lezama and Vice President for Information Systems and Technology (IS&T) Mark Silis for agreeing to act as co-chairs of the Working Group, and to Chief Officer for Business and Digital Transformation Renaud Fournier for leading the Workstream on High-Impact Opportunities.

Background

The Executive Vice President and Treasurer (EVPT) organization provides a wide range of administrative and operational services to MIT – from stewarding the Institute's finances and physical and digital resources, to supporting our workforce, to keeping our community healthy and our campus safe. With the advent of generative artificial intelligence (AI), we are exploring and imagining how AI could be used to enhance efficiency and productivity for campus business operations and enable rapid data-driven decision making and planning.

In fall 2023, MIT appointed a Chief Officer for Business and Digital Transformation, whose team has assembled a seven-year roadmap to modernize our enterprise systems and related business processes. Understanding how AI capabilities can be incorporated into our systems and processes will inform and support this work.

It is within this context that we are convening the **Working Group for AI in Administration and Operations.**

Charge

The Working Group is charged with understanding the fast-evolving landscape of AI and generative AI tools and current capabilities, and how these tools can be applied to make a variety of processes and systems more efficient and effective. These include the tools used for transactional activities across areas such as finance, human resources and auditing. They also include capabilities to control and monitor building operations and occupancy, enhance customer service and address Institute risks by strengthening cybersecurity. The overall goal is to explore opportunities for maximizing the use of these new AI capabilities in ways that significantly improve the experiences of our community and position MIT for the future.

The Working Group will benchmark adoption of AI in business operations at other universities and in industry, and catalog MIT's current use of AI-based tools across the administrative, academic and research areas. It will work to understand what tools, technologies and commercial products are available or in development that could be used at MIT.

The group will deliver a set of recommendations describing opportunities that could be pursued over the next five years. It will also propose a set of core principles to guide MIT's use of these emerging capabilities and identify gaps where we may need to develop policies, guidelines or oversight structures. As part of the process for developing recommendations, the group will initiate or support experiments and pilots, exploring the use of discovery sprints as a methodology to deepen its understanding. The results of the experiments and the

recommendations will be compiled into a roadmap for review and consideration by the EVPT and Provost in collaboration with other MIT senior leaders.

Timeline

It is expected that the group will be convened in September 2024 and submit its final recommendations in the spring of 2025. The co-chairs will keep the sponsors apprised of progress on an ongoing basis.

Sponsors

- Cindy Barnhart, Provost
- Glen Shor, Executive Vice President and Treasurer

Co-chairs

- Armando Solar-Lezama, Distinguished Professor of Computing, MIT Schwarzman College of Computing (SCC); Professor of Electrical Engineering and Computer Science (EECS); Associate Director, Computer Science and Artificial Intelligence Laboratory (CSAIL)
- Mark Silis, Vice President for Information Systems and Technology (IS&T)

Membership

Workstream 1: Setting the Foundation

Members:

- Armando Solar-Lezama, Distinguished Professor of Computing, MIT Schwarzman College of Computing (SCC); Professor of Electrical Engineering and Computer Science (EECS); Associate Director, Computer Science and Artificial Intelligence Laboratory (CSAIL)
- Mark Silis, Vice President for Information Systems and Technology (IS&T)
- Olu Brown, Associate Vice President, IS&T
- Henry Corrigan-Gibbs, Assistant Professor, EECS and CSAIL
- Robin Elices, Executive Director, Office of the Executive Vice President and Treasurer (EVPT)
- Renaud Fournier, Chief Officer for Business and Digital Transformation
- Simon Johnson, Ronald A Kurtz Professor of Entrepreneurship, Sloan
- Deb Roy, Professor, Media Arts and Sciences; Director, MIT Center for Constructive Communication
- Mary Ellen Sinkus, Assistant Provost for Finance

Workstream 2: High-Impact Opportunities

Members:

- Renaud Fournier, Chief Officer for Business and Digital Transformation (Workstream Lead)
- Olu Brown, Associate Vice President, IS&T
- Sean DeBoer, Assistant Dean for Finance and Administration, School of Humanities, Arts, and Social Sciences
- Keith Diggans, Associate Vice President, Facilities Operations, Campus Services & Stewardship
- Robin Elices, Executive Director, Office of the EVPT
- Matthew Eynon, Associate Vice President and Chief Operating Officer, Resource Development
- Andrei Kozhev, Associate Director, Data Analytics, Human Resources
- Colleen Leslie, Director, Assistant Provost for Research Administration, Office of the Vice President for Research (VPR)
- Scott Lever, Associate Vice President, IS&T
- Kathleen McGrath, Director of Financial Operations, Office of the Vice President for Finance (VPF)
- Brian O'Conaill, Director of Administration and Finance, Aeronautics and Astronautics
- Heather Sardis, Associate Director for Technology and Strategic Planning, MIT Libraries
- Brian Schuetz, Executive Director, MIT Health
- Kari Thande, Assistant Dean for Finance, School of Science
- Kate Trimble, Senior Associate Dean of Experiential Learning and Chief of Staff, Office of the Vice Chancellor

Staff Support

Elizabeth McManus, Senior Project Manager, Office of the EVPT

The group may engage consulting assistance to support and expedite this work.

Advisors

Advisors will attend meetings as needed based on the agenda and will be consulted regularly outside of the meetings for guidance.

- Jason Baletsa, Counsel, Office of the General Counsel (OGC)
- Lillian DeWitt, Director of Strategic Sourcing, VPF
- Greg Moffatt, Chief Research Compliance Officer, VPR
- Claire Schneider, Counsel, OGC
- Carter Stubbs, Audit Services Manager, Audit Division

Appendix C: Existing AI Guidance and Tools

1. Published AI Guidance Policy
    A. https://ist.mit.edu/ai-guidance: Information Systems and Technology (IS&T) is providing this initial guidance to encourage community members to consider factors including information security, data privacy, regulatory and policy compliance, compliance with confidentiality restrictions concerning third party information and data, intellectual property (e.g., copyright and patent), and academic integrity when choosing to use or purchase software that makes use of generative artificial intelligence (AI).
2. MIT Advice on AI in Teaching and Learning
    A. https://tll.mit.edu/teaching-resources/course-design/gen-ai-your-course/: To help MIT faculty & instructors navigate the sea of information on the use of generative AI in higher education 2, TLL (Teaching and Learning Lab) has combined timely news, advice, and resources on genAI with best practices in teaching and learning.
3. Written Information Security Program
    A. https://infoprotect.mit.edu/: This new program is based on classifying Institute research data and administrative information according to the risk posed by the loss of confidentiality, integrity, or availability of the information.
4. MIT Licensed AI Products and Tools
    A. https://ist.mit.edu/ai-tools: The products listed here provide or include generative artificial intelligence (AI) tools, services, or components, and are licensed by IS&T for use by the MIT community.
    B. https://ist.mit.edu/ai-faculty (ChatGPT for faculty use)

Generative AI tools available at MIT; available through IS&T for use by MIT

Product/service	Description	How to obtain
Adobe Firefly	Generate or modify images via text prompts	Available to MIT faculty, staff, students, and researchers via MIT's Adobe Creative Cloud license
AWS Bedrock	Developer access to third-party large language models (e.g., Anthropic, A21 Labs, Stable Diffusion) via API	Register for an AWS account through IS&T's MIT Cloud Accounts program
AWS Sagemaker	Managed AI/ML model access and training	Register for an AWS account through IS&T's MIT Cloud Accounts program
Azure OpenAI	Developer access to Azure OpenAI and other cognitive services (e.g., Whisper, DALL-E)	Register for an Azure account through IS&T's MIT Cloud Accounts program
Google Gemini (formerly Bard)	Text-based AI chatbot	Available for MIT Google Workspace accounts
Google Vertex	Developer access to Google's AI models	Register for a GCP account through IS&T's MIT Cloud Accounts program
Microsoft Copilot	Text-based AI chatbot; also capable of generating images	Access Copilot via Touchstone with your MIT Kerberos account
Microsoft Copilot for M365	AI add-on for Microsoft 365 products	Current early access participants only

Product/service	Description	How to obtain
Open AI ChatGPT	Advanced version of publicly available ChatGPT tool with added security and features (for faculty only)	Faculty can request an account (Touchstone authentication required)
Salesforce Einstein	AI add-on for Salesforce	Contact IS&T's Salesforce Licensing
Zoom AI Companion	AI add-on for meeting summaries, AI Companion questions, smart recording, and whiteboard content generation features	Enable features by signing into MIT's Zoom portal via Touchstone

Appendix D: List of Experts

Presentations

- Yoon Kim, Assistant Professor, EECS/CSAIL, MIT: Presentation on "Pitfalls of LLMs/Generative AI" delivered on November 8, 2024. Focused on risks such as hallucinations, language biases, and model instability, with recommendations for human oversight and maintenance.

- Olu Brown, Associate Vice President, IS&T: Presentation on "MIT Admin ChatBot: Beta" delivered on November 1, 2024. Covered technical implementation, data integration, and privacy considerations.

- Klara Jelinkova, CIO, Harvard: Presentation on "Strategic Insights from Peer Institutions" delivered on November 22, 2024. Emphasized hybrid governance, proactive experimentation, and risk mitigation.

- Jen Stave, Executive Director of Digital, Data, and Design, Harvard: Presentation on "Generative AI and Workplace Productivity" delivered on December 6, 2024. Discussed productivity impacts, emotional distress from poor training, and selective investment in AI.

- Deb Roy, Professor, Media Arts and Sciences; Director, MIT Center for Constructive Communication: Presentation on "Constructive Communication and AI Integration" delivered on December 6, 2024. Explored dialogue facilitation, inclusivity, and scaling tools for contentious issues.

Additional MIT Outreach:

- Brian Canavan, MIT Registrar
- Carol Wood, Senior Director, Research Systems, VPR
- Elena Glatman, Executive Director, CSAIL
- Juliana Quattrocchi, Policy and Communications Fellow, Sloan

Appendix E: Benchmarking Insights (as of 2/14/25)

The following provides a high-level summary of peer institute and school benchmarking findings:

Peer School/Institute A

- Leading in AI Strategy with comprehensive website and detailed committee report
- Developing governance through ethics portal (staff section still "coming soon")
- Leading in Enterprise Tools offering multiple platforms (ChatGPT, Copilot, Gemini, LumenGPT)
- Leading in Training with resources for all stakeholders (faculty, staff, students)
- N/A for Pilots & Use Cases as no information is publicly available

Peer School/Institute B

- Leading in Strategy with dedicated AI website and thorough administrative task force report
- Developing Governance through baseline AI guidelines on IT website
- Developing Enterprise Tools centered primarily around Microsoft Copilot
- Developing Training with basic offerings like "How to Talk to an AI Tool"
- Leading in Pilots with organized implementation timeline and a dedicated day for presenting each pilot's success, failures, and next steps

Peer School/Institute C

- Leading in Strategy with robust IT website presence and dedicated taskforce
- Leading in Governance with comprehensive guidelines including specific use scenarios
- Leading in Enterprise Tools offering multiple platforms and secure sandbox environment
- Developing Training focused primarily on introductory/getting started guides
- Leading in Pilots through dedicated Generative AI Administration and Operations Group

Peer School/Institute D

- Leading in Strategy with dedicated website and robust taskforce report
- Leading in Governance with policies integrated into data security framework
- Leading in Enterprise Tools including school-specific chatbot plus Microsoft/Amazon options
- Developing Training through available courses on chatbots and prompts
- Developing Pilots as indicated in taskforce report's forward-looking recommendations

Peer School/Institute E

- Leading in Strategy with dedicated AI website and robust advisory council report
- Leading in Governance through detailed responsible AI guidelines covering safety and risks
- N/A for Enterprise Tools (visible options require SSO access)

- Leading in Training through multiple platforms (LinkedIn Learning, Coursera) and instructor-led options
- N/A for Pilots beyond a robust ‘learn by doing’ sandbox

Peer School/Institute F

- Developing Strategy with basic AI website through IT department
- Leading in Governance with general guidelines, including procurement process
- Leading in Enterprise Tools with dedicated ChatGPT for sensitive/protected data
- N/A for Training with no programs mentioned
- N/A for Pilots with no specific initiatives described

Appendix F: AI Working Group 2 – High-Impact Opportunities

Submitted Use Cases

A total of 41 use cases were evaluated, categorized as follows:

- 4 Resource Development
- 10 Human Resources
- 10 Research & TLO
- 6 Finance
- 4 Student Services
- 7 General Administrative

Evaluation Criteria

Each use case was assessed based on:

- Alignment with the goal of reducing administrative burden
- Feasibility for near-term implementation
- Integration potential with existing AI initiatives (e.g., MIT Admin Chatbot)
- Avoidance of duplication with Business and Digital Transformation projects
- Broad impact across MIT's community

Outcome

9 use cases were identified for their broad benefits and feasibility. These were consolidated into 3 focused pilot projects:

- MIT Chatbot Expansion – Enhancing existing AI chatbot capabilities
- ServiceNow AI Queries & Ticket Submission – Automating service interactions
- Proposal Writing & Federal Solicitation Analysis – AI-driven support for research administration

Next Steps

- Pilots are expected to launch in June 2025, running for 3–6 months
- Approach based on discovery sprints
- Results of the pilots will inform a recommendation report for Glen Shor, EVPT