Amherst College

Task Force on Guidelines for the Use of Generative AI Tools for Teaching and Learning at Amherst

Report

4 December 2023

Scott Alfeld, Computer Science
Riley Caldwell-O-Keefe, Center for Teaching and Learning
Pawan Dhingra, Associate Provost and Associate Dean of Faculty, American Studies
Carmen Granda, Spanish
Jaya Kannan, Academic Technology Services (co-chair)
Chris Kingston, Economics (co-chair)
Kara McGillicuddy, Academic Technology Services
Martha Umphrey, Law, Jurisprudence and Social Thought

Executive Summary

Generative Artificial Intelligence technology ('Gen AI') is likely to dramatically affect the familiar, human-centered paradigm of education offered by liberal arts colleges. This report provides concrete strategies and techniques for negotiating the challenges and opportunities that Gen AI creates.

- **Academic integrity.** The use of AI is covered by the existing Honor Code. Nevertheless, Gen AI poses significant challenges. Our task force noted the need for the College to find ways to articulate broadly and clearly the central values of academic integrity and trust, and to shape the culture and practices around ethical intellectual engagement.

- **Faculty should be free to set policy regarding the use of AI in their courses.** Appropriate strategies will differ across academic disciplines and faculty teaching styles. Whether the goal is to prohibit or embrace the use of AI - or something in between - a variety of approaches are possible. Numerous resources and strategies are provided in this report, and the AmherstAI hub will provide a continuously updated set of resources, including information about AI tools, examples of AI use, and information on upcoming events and educational opportunities.

- **Strategies for teaching and assessment.** Because Gen AI tools can be used to create text, images, video, music, and code that mimic that produced by humans, it may be effectively impossible to enforce prohibitions on the use of AI tools in assignments completed outside of class. Faculty should therefore consider whether and how they may need to adapt pedagogy and assessments in the light of their teaching goals. Assignments that evaluate students based on ***product*** (e.g., writing) completed ***outside class*** may need to be replaced by ***in-class*** assessments or assessments that evaluate the learning ***process***.

- **Transparent and sustained communication.** Whatever approach instructors wish to take to Gen AI, it is vital that faculty clearly articulate their policies and expectations to students both via course syllabi and verbally throughout the semester. Links to resources designed to help faculty design course policies and syllabus statements are available in the report.

- **Other ethical concerns.** Gen AI raises significant concerns about privacy and security of data. Further, Gen AI frequently generates and confidently reports inaccurate, misleading or fabricated content. The content may also reflect biases and stereotypes present in the underlying data.

- **Future directions and governance.** We recommend that a suitable body be charged to review these guidelines periodically and suggest changes in response to evolving technology and experience. As new tools become available, the College will need to develop transparent and consistent policies to govern decisions regarding the adoption and use of new AI tools, with substantial input from faculty. Resources will be required to provide ongoing training, support and guidance to all community members. Finally, the skills that our students will need to thrive in the AI-infused work environments of the future will differ from those needed in the past. The College must consider how our curriculum will need to adapt to these changing needs.

Table of Contents

Query to ChatGPT 3.5, (7 November 2023): How will Generative AI change what it means to be human?

Enhanced Creativity: Generative AI can assist and augment human creativity. It can be used to generate art, music, literature, and other creative works, offering new tools and inspirations for artists and creators. This can expand the boundaries of human creativity and offer new ways to express oneself.

Automation and Labor Market: Generative AI can automate certain tasks and jobs, potentially leading to shifts in the labor market. While it may eliminate some repetitive or dangerous jobs, it can also create new roles related to AI development, maintenance, and oversight.

Augmented Decision-Making: AI can assist humans in making decisions by providing data-driven insights and recommendations. It can analyze vast amounts of information quickly, potentially improving decision quality across various domains, from healthcare to finance.

Language Translation and Cross-Cultural Communication: Generative AI can bridge language barriers by providing real-time translation and interpretation. This can foster global communication and collaboration, potentially leading to more interconnected and diverse societies.

Education and Learning: AI-powered tutoring and educational tools can personalize learning experiences, making education more accessible and effective. This can empower individuals to acquire new skills and knowledge more efficiently.

Ethical and Privacy Concerns: The use of generative AI also raises ethical concerns related to privacy, misinformation, and bias. As AI systems become more sophisticated, addressing these issues becomes increasingly important for preserving human values and rights.

Assistance in Healthcare: Generative AI can help in medical diagnosis, drug discovery, and treatment planning. It has the potential to save lives by providing more accurate and timely medical insights.

Personalization: AI can personalize user experiences in various applications, from content recommendations to healthcare plans, enhancing convenience and relevance in people's lives.

Changes in Communication: AI-driven chatbots and virtual assistants can transform how people interact with technology and each other, blurring the lines between human and machine communication.

Cultural and Social Impact: Generative AI can shape cultural norms and values through the content it generates and the way it influences online discussions. It can be a force for positive change or contribute to divisive or harmful content.

Questions of Identity and Authenticity: As AI systems become more capable of generating text, images, and even videos, it can be challenging to discern what is truly created by humans, raising questions about authenticity and trust in the digital age.

1. Introduction

As the above "self"-description makes clear, Generative Artificial Intelligence technology ('Gen AI') is likely dramatically to affect the familiar, human-centered paradigm of education offered by liberal arts colleges. A relatively new technology that, among other capacities, enables natural-language interaction with computers and can create high-quality text, image, audio, and other content, Gen AI promises to transfigure and hybridize the relationship between human and computer around which we have organized our educational model over the past few decades. (A more detailed resource describing the current capabilities and limitations of the technology can be found here).

Gen AI will inflect our teaching and learning environment in ways reminiscent of the introduction of the internet. On the one hand, it opens exciting new possibilities to develop certain capabilities: AI tools can enhance students' problem-solving abilities and critical thinking skills, and expand their interdisciplinary knowledge. Amherst students already use Gen AI tools as aids in language courses, in researching and summarizing articles, in brainstorming ideas and generating questions, in creating visualizations, and the like. AI can help instructors adapt teaching tools to the capabilities and learning styles of individual students, enhancing our ability to provide individualized learning experiences that meet the needs of students who may demonstrate learning in multiple ways.

On the other hand, Gen AI's power to create and respond (not just find, collect, and synthesize) raises questions about many of our fundamental, long-held educational presuppositions. Why, for example, should students learn the fundamentals of reading, writing, and creating if a machine can summarize and compose? Might the availability of these tools interfere with and inhibit intellectual growth by relieving students of the need to master basic skills? How might Gen AI augment class materials (via correction, summary, translation, and individualized calibration) outside the student-faculty relationship, and with what effects? How should we address serious concerns about privacy and the perpetuation of harmful biases in AI-generated content? What values and skill sets (intuiting, close reading, speaking) does a liberal arts education impart that cannot be generated or replicated by technology, and how should our curriculum respond?

Because Gen AI technology is evolving at an extremely rapid pace, the Task Force believes it is incumbent upon the Amherst community to articulate a responsible, values-based approach to its adoption. Indeed, time is of the essence: students, faculty, and staff have already begun to use Gen AI tools in their work, and the technology is rapidly becoming more powerful.

The Amherst College Generative AI Task Force was charged with developing guidelines to provide to the community on **classroom and pedagogy related topics**, including best practices for integrating AI tools into teaching; instructors' responsibilities to clearly communicate expectations regarding AI use to students; implications for academic integrity; attention to issues of equity and accessibility; and strategies to keep the Amherst community informed. Our work benefited from conversations with Director of Community Standards Corey Michalos, Director of the Strategic Learning Center Larissa Hopkins, and Chief Information Officer David Hamilton, and from feedback from staff and faculty participants as well as a student panel who shared their experiences as part of the Fall 2023 AI Learning Lab lunch series organized by ATS.

In what follows we offer an overview of the challenges and opportunities we have identified in our discussions about Gen AI over the Fall 2023 semester. We also provide concrete strategies and techniques for negotiating the challenges and opportunities that Gen AI creates and incorporating Gen AI into our pedagogy. Much work remains, work that must continue and evolve as these technologies evolve. The Gen AI Task Force recommends that future efforts – whether done via committee decisions, curricular development, or administrative action – be guided by the following precepts:

- The College must **reaffirm the centrality of human relationships** as we work to foster intellectual engagement and development, and affirm to our students that they possess the ability to generate original ideas, that their ideas hold significance, and that education involves more than the generation of content.

- **Faculty must remain free to determine how to adapt their teaching** to achieve their learning goals and promote ethical academic engagement in the presence of AI, and to decide when and how to use these tools, if at all, in their teaching.

- The College must **find ways to articulate broadly and clearly the central values - indeed the very meaning - of academic integrity and trust**, and to address the culture and practices around ethical intellectual engagement, which are being tested by the introduction of GenAI into our learning environment.

- The College must **develop appropriate capacities to support ongoing education of faculty, staff and students** to meet the fast-changing Gen AI environment, to understand the vulnerabilities and biases baked into its informational model, and to articulate a vision of the values and skills students will need to succeed in a Gen AI-infused world.

- The Faculty, whose responsibility is to oversee the college's curriculum, must be involved in **decisions about the adoption and uses of Gen AI technologies that affect academic work** both across the college and in individual classrooms.

Our goal in this report is to provide guidelines and practical tools to help faculty decide the extent of AI use that will be most compatible with achieving their teaching/learning goals (on a spectrum from an AI-free to an AI-rich environment), and to suggest strategies for faculty who may take a variety of positions on AI use to craft approaches with pedagogic purpose in order to realize its potential benefits while mitigating its negative impacts.

We noted and discussed the possibility that broader curricular changes may be desirable and perhaps necessary. Future work environments will include AI tools. The skills that our students will need to thrive in such environments will differ from those that they have needed in the past. As a college, we will need to consider how the curriculum should respond to those demands. For example, the skills required to write successfully in an AI world may differ from those needed previously, just as handwriting and long division have been displaced by keyboards and calculators, and memorization of facts has been displaced by Google search engines. How should our teaching evolve to meet these needs, build digital and non-digital literacy, and maintain the relevance of a liberal arts degree? Should AI literacy be intentionally introduced into the curriculum? Where?

Though beyond the scope of our charge, we noted also that the College must find ways and constitute structures that allow us to think ahead, collectively, about the coherence of stem-to-stern institutional changes that Gen AI will both demand and enable – not just in teaching but in research, admissions, student affairs, career advising, athletics, operations, alumni relations, and beyond. The perspectives of faculty, students and staff are all important to include in these conversations and should be part of process and policy decisions.

2. Ethical concerns

Generative AI has given rise to a variety of ethical concerns, including effects on employment, environmental impact, and the potential for misinformation (including “deepfake” images and videos). Consistent with the limited scope of our charge, here we focus on three issues that are particularly salient in our context: Academic Integrity, Privacy and Security, and Bias.

2.1 Promoting Ethical Academic Engagement

The use of AI falls under the broader academic integrity policy and is covered by the existing Honor Code. However, the technology poses significant challenges that may require us to adjust our practices if we are to maintain ethical academic standards. Our task force noted the need for a general cultural shift at Amherst addressing core issues of cheating, trust, and integrity. It is important that space for dialogue about these core values be created either

during orientation, first-year seminars, in our own courses, or through some other ongoing programmatic effort.

> **Statement of Intellectual Responsibility**
>
> Every person's education is the product of their intellectual effort and participation in a process of critical exchange. Amherst College cannot educate those who are unwilling to submit their own work and ideas to critical assessment. Nor can it tolerate those who interfere with the participation of others in the critical process. Therefore, the College considers it a violation of the requirements of intellectual responsibility to submit work that is not one's own or otherwise to subvert the conditions under which academic work is performed by oneself or by others.

Amherst's Statement of Intellectual Responsibility considers it a violation of intellectual responsibility "to submit work that is not one's own." But in the age of generative AI, what does it mean to submit "one's own" work? While students often seek help from, for example, the Writing Center and Moss Quantitative Center as they draft essays or work on problem sets, programs such as ChatGPT, Quivr, and others can generate, stylize, critique, and edit writing at a level already adept enough to earn decent grades at top universities. What lines can and should be drawn to differentiate between writing that is one's own and writing co-authored with technology?

Arguably, Gen AI can be said to "subvert the conditions under which academic work is performed by oneself." Students can use ChatGPT, for example, to brainstorm, obtain information, outline, draft, critique, and redraft essays. Those uses cannot at present be policed. Faculty must be free to set course policy regarding the use of AI in their courses; but whatever their approach, faculty may need to adapt assessments and consider how to clearly communicate their expectations around the ethical use of AI to students. As Gen AI's capabilities expand, our understanding of individual authorship may necessarily need to shift to accommodate the role it plays in student work.

Such a shift may require us to revisit the language in our current Statement of Intellectual Responsibility. At present, in order to make transparent our classroom expectations around academic integrity, it is vital that faculty include in their syllabi a clear statement concerning the use, if any, of Gen AI technologies in individual courses, in order to help students navigate the ethical use of AI in this changing environment.

All faculty should be aware that Amherst's Office of Community Standards relies on syllabus statements in individual courses to establish standards of conduct when asked to hold students accountable for academic integrity violations.

- **Strategies for Teaching and Assessment.** Faculty should consider whether and how to adapt assessments in the light of their teaching goals. Whether the goal is to prohibit or

embrace the use of AI - or something in between - a variety of approaches are possible. Resources and examples are provided in section 3 below.

- **Syllabus Statements.** Course syllabi should contain a clear statement of the ways in which generative AI tools may or may not be used by students in a course and on specific categories of assignments. Instructors should communicate these expectations and their rationale verbally to students at the start of and throughout the semester. Students should be encouraged to ask for clarification as needed. Syllabi should also provide guidance to students about how they are expected to document and attribute the use of AI tools in their work, and to validate or verify output produced using AI. Sample syllabus statements, and links to resources designed to help faculty design course policies and syllabus statements, are available in the Pedagogy Resource Guide available via the Amherst AI hub.

- **Detection Software.** It is very difficult to accurately detect, much less prove, whether content has been generated by/with AI. The use of detection software (such as Turnitin) is strongly discouraged. Such services have been shown to yield unreliable results, including both false positives and false negatives. Further, the results are likely to be biased, for example, against non-native English writers. Moreover, uploading student-produced content to AI detection software without their prior permission may breach student privacy and intellectual property.

- **Department conversations.** Departments are strongly encouraged to hold ongoing conversations about the use of AI with all faculty (including non-tenured, teaching staff, etc.), and to aim for consistency while respecting faculty rights to set appropriate policies in individual courses. In those conversations, departments might consider adding a default AI policy to their departmental handbook.

2.2 Privacy and Security

Many AI tools do not have robust privacy controls. Confidential data such as research data, salary details, disciplinary information, student transcripts, resumes, job application materials, etc., should not be shared with AI tools. Doing so could expose sensitive information to unauthorized parties. It is the instructor's responsibility to safeguard student data following all relevant regulations covered by FERPA (Family Educational Rights and Privacy Act). If you are considering licensing a new tool or utilizing a new AI feature in a tool you already use, **please work with AskIT** to ensure that the tools and services you procure or are using on behalf of the College have appropriate privacy and security protections and are assessed for risk prior to use. If you have already begun the use of new tools or features, configure the platform so that it

does not use or retain chat/prompt history, if possible. Once data is placed into these platforms, there are no guarantees that they can later be removed from them.

2.3 Inaccuracy and Bias in AI-generated content

Gen AI frequently generates and confidently reports inaccurate, misleading or fabricated content, sometimes referred to as "**hallucinations**," including plausible-sounding but imaginary references, false assertions of "facts," etc. In addition, Gen AI tools can and do produce content that may reflect **bias and stereotypes** present in the underlying data. A valuable guide to the responsible use of AI tools produced by Frost Library notes that "One way bias occurs is through datasets that misrepresent, exclude, or marginalize certain social identities, communities, and practices. When models are trained on these datasets, they will reflect and often amplify social prejudices and stereotypes."

- Students and faculty should be cautious and verify all AI-generated content with a credible source. Faculty are responsible for verifying the accuracy of any content they produce or publish that includes AI-generated material.
- We encourage faculty and teaching staff to utilize library resources and, in particular, the expertise of our research librarians in educating students about the reliability of sources via, for example, class visits as part of the first-year seminar program and in writing-intensive courses.
- As with the internet more generally, content created by generative AI reproduces assumptions and biases prevalent in society at large. This is all the more reason to check the validity of any content.

3. Teaching Strategies in the Age of Gen AI

Generative AI tools will have substantial impacts on teaching and learning, both enhancing and disrupting our pedagogy in various ways. AI tools are already part of the machinery in many writing and calculation platforms, including Microsoft products and Grammarly, and are rapidly proliferating. Although "chatbot" Gen AI platforms such as ChatGPT generally cannot (so far) respond particularly well to humanities-type prompts, many students have already begun to use these tools as aids in doing their academic work. Instructors should now expect that any work completed outside of class will frequently involve the use of generative AI.

Some faculty have already begun to integrate generative AI into their courses and pedagogical strategies; others find the prospect of it highly disruptive to their courses' objectives and aspirations. Faculty should be free to set course policy regarding the use of AI in their courses. We recognize that appropriate uses will differ across academic disciplines and faculty teaching styles, and that our understanding of these tools will expand as we use them. In responding to these challenges, instructors will need to consider:

- how best to align our learning objectives with the role we would like technology to play in our courses and with the learning needs of students who will enter a world filled with AI tools.
- how to assess student work in order to provide incentives for students to invest in the learning goals and processes we hope to teach them.
- how to communicate course policies and expectations to students.
- how to ensure that AI use promotes accessibility and equity.

In this section, we provide resources that we hope will assist faculty in considering these adjustments. Because the technology is rapidly evolving, we have created a Pedagogy Resource Guide on the Amherst AI hub as a living document which will be continually updated by the Center for Teaching and Learning to provide relevant pedagogical strategies.

Whatever approach instructors wish to take to Gen AI, whether to curtail its use as much as possible or to fully allow and embrace it, it will be important to clearly articulate in our syllabi and throughout the semester what policies we expect our students to follow, including making students aware that (as discussed in section 2) AI tools pose privacy and security risks and are prone to make false statements and exhibit biases. As noted above, sample syllabus statements and related resources are available in the Resource Guide.

Below, we consider three approaches that an instructor might decide to take in any particular course, depending on the extent of AI use that will be most compatible with achieving learning and teaching goals for each course: (3.1) Limiting students' Gen AI use students as much as possible; (3.2) allowing selective use of Gen AI; and (3.3) embracing students' use of AI as a key tool for learning.

3.1 Strategies to limit Gen AI usage as much as is practicable.

Because Gen AI tools can be used to create text, images, video, music, and code that mimics that produced by humans, it may be effectively impossible to enforce prohibitions on use of AI tools in assignments completed outside of class. Consequently, assignments that evaluate

students based on ***output*** (e.g., writing) completed ***outside class*** may need to be replaced by in-class assessments or assessments that evaluate the learning process rather than product. Imaginative changes in assessment methods can reduce students' incentives and ability to cheat, while also enhancing their learning. This resource from the University of Michigan provides instructors with a set of reflective questions to determine whether or not their assignments need revisions when considering teaching and learning in a Gen AI era. Suggestions for modes of assessment that minimize the impact of AI include:

- In-class, especially handwritten, assessments (e.g., in-class tests or written responses completed in class or during scheduled exams).
- In-class discussions or other oral assignments such as role-playing, debates, or interactive presentations.
- Assignments that require students to critically engage materials unlikely to be accessible to Gen AI tools. These might include primary materials (e.g., material in the Mead or the College Archives, or interviews) and class activities (lecture, speakers, field work, lab, discussion).
- Assignments or assessments that must be completed in stages (e.g., proposal, outline, rough draft, and final draft) and on which they receive instructor and peer feedback throughout the process. This video from Harvard has a useful discussion with concrete ideas for rendering process (rather than product) visible in writing courses.
- Incorporate metacognitive elements: asking students to reflect on their process, decision-making, how their understanding evolved over the course of their research, what they learned from the exercise, where they hit roadblocks or achieved breakthroughs, and how they would approach it differently in the future.
- Require students to enable Google's "revision history" capabilities to allow instructors to review the iterative process of their work. Note: it is possible to create Google assignments in Moodle to enable instructors to easily see the revision history.
- Use an annotation tool such as Perusall to ask students to engage in annotating their work or course materials.
- Explore how a Gen AI tool responds to assignment prompts. If the results are quite strong, consider revising the prompt to strengthen the complexity or specificity of problems students are engaging. AI is less adept at critical thinking skills such as comparing phenomena than at summarizing information.

3.2 Strategies for selective use of Gen AI tools

Some instructors may decide that the use of AI can enhance learning in some parts of the course, but be detrimental in other parts, and therefore wish to explicitly allow or encourage AI use in some parts of the course while prohibiting it in others. A course might, for example, combine in-class exams that exclude AI use with take-home projects or essays for which some kinds of AI use are allowed, in order to ensure both that students gain foundational skills and learn to use AI to enhance their productivity in preparation for real-world challenges.

In this kind of scenario, it is vital to convey the permissible uses of AI in a course clearly to students, both in the syllabus and in person. Students should be given clear guidelines as to how they are expected to acknowledge and document their AI use, and be encouraged to ask questions about permissible AI use and to check all AI outputs for accuracy.

- This article from the University of Sydney discusses achieving balance between assessments that use AI and those that do not, and links to a guide to revising assessments in ways that ensure the effective, transparent and ethical use of AI (here is a related webinar).
- This tool from GWU will help faculty decide and articulate to students their decisions around how GenAI may be used in various aspects of the course.
- Harvard MetaLab's AI Pedagogy Project provides a "collection of curated assignments that integrate AI tools" that may help faculty generate ideas for their own courses.
- In this video from Harvard's Bok Center for Teaching and Learning, faculty discuss strategies and provide examples of adapting assignments for a writing course.

3.3 Strategies for incorporating Gen AI as a key tool for learning

Some instructors may want to incorporate AI tools into their assignments in order to help students build the skills necessary to use these tools while remaining mindful of their limitations (including privacy concerns, inaccuracy, and biases as discussed in section 2 above). Students might potentially use Gen AI tools for all stages of idea development, creation, and revision of a project. They might also, for example, be asked to keep a journal throughout the semester that reflects critically on their use of AI and its impact on their project's development. Instructors can teach students "prompt engineering" strategies to maximize the quality of generative AI responses. As emphasized above, it will be important that the rationale for AI use and the rules governing its use, including citation practices, be transparently communicated to students both via syllabus statements and verbally in class.

There are many possible uses of Gen AI to enhance teaching, and more ideas will undoubtedly emerge as the technology matures and instructors explore its capabilities. As noted above, a Pedagogy Resource Guide (available via the Amherst AI hub) will maintain a list of curated examples. Here are more examples from Yale, Harvard, and Amherst. A few illustrative examples include:

- Diffit: A tool that tailors materials to match students' individual levels of knowledge. It can be used, for example, to enhance language instruction by generating versions of a text suitable for learners at different proficiency levels.
- For students studying languages, this resource provides ideas for using AI as an interlocutor for role-playing exercises.
- Use of ChatGPT as a tool to explore the use of metaphor in a poetry class (Harvard).
- Use of Gen AI as a virtual "teaching assistant" that can help instructors devise explanations, prompts, or scenarios that make sense to their students, and creating low-stakes quizzes. Magic School can help create various kinds of course elements, such as drafts of customized rubrics for various assignments, that can then be honed by the instructor.
- An interesting one-hour information session from Harvard discusses ideas for using AI to enhance STEM teaching.

4. Future directions: governance, communication, resources and curriculum

The charge to this Task Force focuses on classroom teaching. We believe the College must also consider how AI will impact the work of all members of the community, including staff. In this section, we offer brief recommendations for institutional changes to help the College facilitate this conversation and respond to future challenges and opportunities. Our report comes at a time of rapid change and growth in the AI industry and widespread speculation and fear about its future consequences. Given this uncertainty, our recommendations are just one contribution to a community-wide conversation that has already begun and will only grow in importance.

- **Future governance.** We recommend that the Faculty Executive Committee consider revising the charge and makeup of the Faculty Computer Committee (FCC), or some

other suitable body, to review these guidelines periodically and make further recommendations in response to evolving technology and experience.

- **Rules/procedures for adoption of AI tools.** As new tools become available, the College will need to develop transparent and consistent policies to govern decisions regarding institutional adoption of AI tools (such as the recent decision to provide Grammarly). These are difficult questions that potentially impinge on faculty's pedagogical autonomy and staff expertise. In some cases, AI tools may enhance accessibility for students with particular learning needs, disabilities, etc. Further, deciding against providing access to tools may disadvantage students who lack the means to purchase individual access. At the same time, there may be good reasons not to adopt tools that affect our ability to teach the values and skills that we want our students to learn. For example, automating the process of generating ideas and editing writing may deprive students of opportunities to learn how to craft cogent and original arguments in their own authentic voices, rather than predictable arguments using generic language. It is therefore important that a transparent policy with clear principles and mechanisms for adoption decisions be developed, involving substantial faculty input for decisions regarding tools that affect pedagogy. Such faculty consultation and oversight could, for example, come from a revamped Faculty Computing Committee or from the Committee on Educational Policy.

- **Curated online resources.** The AmherstAI hub, developed by Academic Technology Services, together with the Center for Teaching and Learning, will provide a continuously updated set of valuable curated resources, including pedagogical resources, information about AI tools, examples of AI use, guidelines for AI use, this report, and information on upcoming events and educational opportunities.

- **Communication.** We recommend that the relevant faculty governance body (e.g., a revamped Faculty Computer Committee) provide regular updates to faculty by email and at Faculty Meetings, and develop mechanisms to gather feedback from students, faculty and staff to share experiences, identify areas of concern, opportunities to use AI to enhance teaching, and best practices, in order to catalyze broader institution-wide discussions involving the whole community.

- **Provost's Retreat.** We suggest that an upcoming annual Provost's Retreat be devoted to discussion of the impact of AI on the liberal arts, including a discussion of assessment strategies, examples of the use of AI to enhance teaching, and potential curricular innovations.

- **Training and ongoing faculty development**. With guidance and support from the Provost's Office, Faculty Computer Committee, and IT, ATS/CTL should organize ongoing training for faculty interested in gaining hands-on assistance with adapting courses to AI. Via the AmherstAI hub, ATS/CTL will also play an important role in facilitating the sharing of knowledge and ideas among faculty and instructional staff.

- **Resources**. IT will likely need staff conversant in AI to train and support faculty, students and instructional staff in navigating their interactions with AI and in providing students with clear guidance on how to use generative AI to support their learning, assignments, and research. The College must also commit to providing the required infrastructural support, including funding and professional development opportunities, for the staff required to carry out these training and support functions.

- **Curriculum**. The courses we teach, and their content, continually evolve in response to shifts in technology, development of new kinds of knowledge and ways of knowing, student interest, and other factors. The skills that our students will need to thrive in the AI-infused work environments of the future will differ from those needed in the past. Accordingly, our curriculum will need to adapt to these changing needs, building both digital and non-digital literacy, if we are to preserve the value and relevance of a liberal arts degree. We recommend that the relevant committees, or perhaps a curricular task force that includes students, take up the question of whether and how AI literacy (e.g., prompt engineering training) should be intentionally introduced into the curriculum, and whether this curricular evolution should be allowed to evolve organically or be consciously coordinated and directed.

Even as generative AI poses significant challenges to some of our traditional educational practices, its emergence calls on us to answer far-reaching questions about the very purpose of a liberal arts education. In that light, we urge the college to imagine and develop new curricular initiatives that cultivate the skills and values we believe are essential to, though not necessarily generated by, the AI-infused world we now inhabit: public speaking and group collaboration skills, perhaps; but also, fundamentally, the values of integrity and the pursuit of truth in an era of disinformation, of deliberativeness and attentiveness to process in a world that overemphasizes output, of responsible experimentation, of personal expression, and of intellectual courage and adaptability. We believe that Amherst is precisely the place to sow the seeds of a confident, ethical vision of human capacity and worth – an absolutely vital project in this new technological age.

Appendix: Charge to the Task Force on Guidelines for the Use of Generative AI Tools for Teaching and Learning at Amherst

Generative AI tools promise to continue to result in opportunities, as well as challenges, for teaching and learning at institutions of higher learning. The pace at which these technologies are evolving suggests that having guidelines in place for their use would be helpful for many within the Amherst community—contributing to greater understanding about the implications of engagement with these tools in an educational context, and serving as a source of best practices for faculty and students. At the same time, it is clear that any guidelines that are developed will themselves need to evolve over time—as the terrain of AI continues to shift, and navigation must be adjusted accordingly.

The task force is charged with developing guidelines that will address the areas listed below, as well as any others that the members feel would be useful. In undertaking this work, the task force should consult with students, faculty, and staff, as needed, and, potentially, with colleagues at peer institutions that have already put guidelines/best practices in place, as well as drawing on other external resources. In creating guidelines, providing specifics will be particularly welcome.

- Approaches/best practices for integrating AI tools into teaching (e.g., assignments, readings, reports, feedback to students, assessment)
- Responsibilities of the instructor to make expectations clear to students about the permissible use of AI tools for coursework and about course-specific policies that instructors may wish to use.
- The implications of AI tools in the area of academic integrity and recommendations on whether updates to current policies are needed and, if so, what these updates would be. Related issues to address might include guidance about the verification of sources, attribution, documentation, and ethics.
- The faculty's responsibilities when publishing content that includes AI-generated material in course materials.
- Attention to issues relating to equity and accessibility
- Strategies to inform the Amherst community about the use of AI tools and to generate discussion that will highlight salient issues.

These are large and complex questions. Nevertheless, it is essential that the task force develop draft guidelines by December 4, 2023, so that some guidelines can be put in place expeditiously. Once completed, the document will be forwarded to the Faculty Executive Committee.