



# Making AI Generative for Higher Education

Adoption and Challenges Among Instructors and  
Researchers

Claire Baytas  
Dylan Ruediger

May 1, 2025



ITHAKA S+R

Ithaka S+R provides research and strategic guidance to help the academic and cultural communities serve the public good and navigate economic, demographic, and technological change.

Ithaka S+R is part of ITHAKA, a nonprofit with a mission to improve access to knowledge and education for people around the world. We believe education is key to the wellbeing of individuals and society, and we work to make it more effective and affordable.

Copyright 2025 ITHAKA. This work is licensed under a Creative Commons Attribution Non-commercial 4.0 International License. To view a copy of the license, please see <https://creativecommons.org/licenses/by-nc/4.0/>.



# Introduction

Generative artificial intelligence (AI) has been a buzz word across higher education ever since OpenAI announced the commercial release of ChatGPT in November 2022. Two and half years later, determining how generative AI will impact and is already impacting teaching, learning, and research—as well as what types of governance need to be put into place to manage that impact—remain priority issues for stakeholders across the sector.

In the immediate wake of ChatGPT's release, student academic integrity was top of mind: the difficulty in detecting content generated by artificial intelligence led instructors to question how their previous plagiarism policies for student work could still be enforced.<sup>1</sup> As time has gone on, the conversations around generative AI have become more nuanced. Stakeholders across higher education have been actively exploring whether and how the technology can enhance teaching, learning, and research. Discussions have also focused on the ethical and societal impacts of the technology, especially the risks related to data security, inaccuracy, and bias.<sup>2</sup> Meanwhile, major technology companies have

---

<sup>1</sup> Examples of discussion around student academic integrity include: Collin Binkley, "Cheating on College Essays? Some Use ChatGPT to Write Them," *AP News*, January 16, 2023, <https://apnews.com/article/chatgpt-cheating-ai-college-1b654b44de2d0dfa4e50bf0186137fc1>; Tom Muir, "Will ChatGPT Change Our Definitions of Cheating?" *Times Higher Education*, November 2, 2023, <https://www.timeshighereducation.com/campus/will-chatgpt-change-our-definitions-cheating>; "The Latest Insights Into Academic Integrity: Instructor and Student Experiences, Attitudes, and The Impact of AI 2024 Update," *Wiley*, July 26, 2024, [https://www.wiley.com/en-us/network/education/instructors/teaching-strategies/the-latest-insights-into-academic-integrity-instructor-and-student-experiences-attitudes-and-the-impact-of-ai-2024-update?utm\\_medium=pressrelease&utm\\_source=wileynewsroom&utm\\_content=augustpressrelease&utm\\_term=academicintegrityreport](https://www.wiley.com/en-us/network/education/instructors/teaching-strategies/the-latest-insights-into-academic-integrity-instructor-and-student-experiences-attitudes-and-the-impact-of-ai-2024-update?utm_medium=pressrelease&utm_source=wileynewsroom&utm_content=augustpressrelease&utm_term=academicintegrityreport).

<sup>2</sup> Examples of recent analyses of the ethical and societal impacts of AI include: Ming Li, Ariunaa Enkhtur, Beverley Anne Yamamoto, Fei Cheng, and Lilan Chen, "Potential Societal Biases of ChatGPT in Higher Education: A Scoping Review," *Open Praxis* 17, no. 1 (2025): 79–94, <https://doi.org/10.55982/openpraxis.17.1.750>; Faye-Marie Vassel, Evan Shieh, Cassidy R. Sugimoto, and Thema Monroe-White, "The Psychosocial Impacts of Generative AI Harms," *Proceedings of the AAAI Symposium Series* 3, no. 1 (2024): 440–47, <https://doi.org/10.1609/aaaiss.v3i1.31251>; Evan Shieh, Faye-Marie Vassel, Cassidy R. Sugimoto, and Thema Monroe-White, "Laissez-Faire Harms: Algorithmic Biases in Generative Language Models," *arXiv preprint*, last revised April 16, 2024,

continued churning out new versions of large language models, while vendors have introduced new product features. For the higher education market specifically, the landscape of generative AI products is sizable and growing, with tools for researchers, teachers, and students that assist with discovering and understanding information, generating and revising writing, and more.<sup>3</sup>

Higher education institutions have reacted to varying degrees to the advent of generative AI. Most universities now have AI task forces, provide sample AI policy language for syllabi, and offer workshops on basic AI literacy. Certain campuses have also begun providing generative AI access to their communities: for example, the University of Michigan, Arizona State University, and the California State University system made headlines when they announced partnerships with big tech companies to put LLM-powered tools in the hands of faculty, students, and staff.<sup>4</sup> Academic publishers have been crafting publication guidelines for AI use, while scholarly societies and other organizations in and around higher

---

<https://doi.org/10.48550/arXiv.2404.07475>; Iain Weissburg, Sathvika Anand, Sharon Levy, and Haewon Jeong, “LLMs Are Biased Teachers: Evaluating LLM Bias in Personalized Education,” *arXiv* preprint, last revised February 9, 2025, <https://doi.org/10.48550/arXiv.2410.14012>.

<sup>3</sup> For further analysis of the product landscape, see Claire Baytas and Dylan Ruediger, “Generative AI in Higher Education: The Product Landscape,” *Ithaka S+R*, March 7, 2024, <https://doi.org/10.18665/sr.320394>, and our Generative AI Product Tracker, <https://sr.ithaka.org/our-work/generative-ai-product-tracker/>.

<sup>4</sup> Tom Burns, “ITS Debuts Custom Artificial Intelligence Services Across U-M,” *The University Record*, August 21, 2023, <https://record.umich.edu/articles/its-debuts-customized-ai-services-to-u-m-community/>; Annie Davis, “A New Collaboration with OpenAI Charts the Future of AI in Higher Education,” *ASU News*, January 18, 2024, <https://news.asu.edu/20240118-university-news-new-collaboration-openai-charts-future-ai-higher-education>; “CSU Announces Landmark Initiative to Become Nation’s First and Largest AI-Empowered University System,” *CSU News*, February 4, 2025, <https://www.calstate.edu/csu-system/news/Pages/CSU-AI-Powered-Initiative.aspx>. On the AI computing center for the Empire AI consortium of institutions in the state of New York, see “Governor Hochul Unveils Fifth Proposal of 2024 State of the State: Empire AI Consortium to Make New York the National Leader in AI Research and Innovation,” *Governor Kathy Hochul*, January 8, 2024, <https://www.governor.ny.gov/news/governor-hochul-unveils-fifth-proposal-2024-state-state-empire-ai-consortium-make-new-york>. On the collaboration between Google and the Stephen M. Ross School of Business at the University of Michigan to launch a Virtual Teaching Assistant pilot program leveraging agentic AI, see “Google Public Sector Helps Enhance Learning at the University of Michigan with Pioneering New Agentic AI Virtual Teaching Assistant,” *Michigan Ross News*, April 7, 2025, <https://michiganross.umich.edu/news/google-public-sector-helps-enhance-learning-university-michigan-pioneering-new-agentic-ai>.

education have also assembled task forces and developed support resources.<sup>5</sup>

Universities recognize the need to coordinate institution-wide support for crucial AI initiatives, such as fostering AI literacy among students, faculty, and staff.

Amidst this flurry of activity, unresolved questions remain when it comes to generative AI's integration into postsecondary teaching, learning, and research. Universities recognize the need to coordinate institution-wide support for crucial AI initiatives, such as fostering AI literacy among students, faculty, and staff.<sup>6</sup> However, institutional silos and decentralized decision-making processes make achieving this goal difficult. The financial implications of going all in on AI for academic institutions remains unclear.<sup>7</sup> Managing student academic integrity policies is still a challenge, and many feel that publisher policies for researchers should be more robust as well.<sup>8</sup> Scholarly inquiry into best practices for integrating AI into teaching, learning, and research is proliferating, but has inevitably struggled to keep up with the pace at which these new technologies are being put in the hands of the community, meaning that many have been learning on the fly.

---

<sup>5</sup> For an example of publisher guidelines, see Wiley's recent release of guidelines for responsible and effective use of AI in authorship: "Wiley Releases AI Guidelines for Authors," *Wiley Newsroom*, March 13, 2025, <https://newsroom.wiley.com/press-releases/press-release-details/2025/Wiley-Releases-AI-Guidelines-for-Authors/default.aspx>. For examples of task forces and resources on AI organized by scholarly communities, see the MLA-CCCC joint task force (<https://aiandwriting.hcommons.org/>) or the ARL-CNI Joint Task Force on Scenario Planning for AI/ML Futures (<https://aiandwriting.hcommons.org/>).

<sup>6</sup> Ithaka S+R's newly launched cohort project focuses on helping universities manage the challenge of integrating AI literacy into curricula; see Ruby MacDougall, Dylan Ruediger, Nathan Kelber, and Zhuo Chen, "Integrating AI Literacy into the Curricula," *Ithaka S+R*, April 9, 2025, <https://sr.ithaka.org/blog/integrating-ai-literacy-into-the-curricula/>.

<sup>7</sup> For a discussion of how universities have implemented generative AI systems, including the financial implications, see: Coalition for Networked Information, "Research University Strategies for Implementing Generative Artificial Intelligence systems," *YouTube*, November 26, 2024, 49:57, [https://www.youtube.com/watch?v=v-E3Mn\\_iHLU](https://www.youtube.com/watch?v=v-E3Mn_iHLU).

<sup>8</sup> For a recent analysis of how publisher policies could be more robust, see Avi Staiman, "When Declarations Just Don't Cut It: Building a Risk-Based Framework for AI Guidelines in Publishing," *Science Editor* 48 (2025), <https://doi.org/10.36591/SE-4801-05>.

In fall 2023, Ithaka S+R launched a collective research project with the objective of studying generative AI's impact on teaching, learning, and research at the postsecondary level.<sup>9</sup> Through a collaboration with 19 universities from across the US and Canada, the “Making AI Generative for Higher Education” project has provided an opportunity for co-learning among cohort members, gathering and sharing data about instructors’ and researchers’ practices when it comes to generative AI, and for leveraging design-thinking to envision new forms of AI-related support. The project also led Ithaka S+R to explore the generative AI product landscape and launch its tracker of generative AI products for higher education.<sup>10</sup>

The present report presents the findings of the interviews conducted by Ithaka S+R and teams from our 19 cohort institutions during the spring of 2024.<sup>11</sup> These interviews asked faculty, graduate students, and other individuals to reflect on their perceptions of and experiences with generative AI in both teaching and research contexts. While the full interview guide can be found in Appendix C, our study was driven by the following questions: To what degree are instructors and researchers adopting generative AI, and how is this changing their approaches and practices in teaching and research? What challenges are they facing in the aftermath of generative AI's emergence? What kinds of support have they benefited from, and what kinds of support do they still need?

---

<sup>9</sup> Danielle Miriam Cooper and Dylan Ruediger, “Making AI Generative for Higher Education: Announcing the Partners for a New Multi-Year Research Project,” *Ithaka S+R*, May 24, 2023, <https://sr.ithaka.org/blog/making-ai-generative-for-higher-education-2/>.

<sup>10</sup> Claire Baytas and Dylan Ruediger, “Generative AI in Higher Education: The Product Landscape” *Ithaka S+R*, March 7, 2024, <https://doi.org/10.18665/sr.320394>. Ithaka S+R’s Generative AI Product Tracker, <https://sr.ithaka.org/our-work/generative-ai-product-tracker/>.

<sup>11</sup> The report also supplements interview findings with other research in the field from the past year, to highlight parallels in findings as well as foreground concerns that still remain pressing for the higher education sector today.

# Key Findings

- Instructors and researchers have widely varied degrees of familiarity with AI, but even those at the lower end of the scale recognize the importance of improving their AI literacy levels.
- Instructors are taking it upon themselves to integrate basic AI skills into student activities but are still determining how generative AI can help them meet course learning objectives and how/if to reimagine those learning objectives.
  - Instructors desire further top-down guidance related to student academic integrity and the formal integration of AI literacy into student general education.
- Most researchers have already experimented with AI, but far fewer have settled on productive ways of integrating the tools for the longer term.
  - Researchers seek further clarity around ethical standards and best practices to ensure research quality and integrity can be maintained.
- Instructors and researchers see a gap in discipline-specific support resources at their institutions and are concerned about having secure, affordable access to generative AI tools. They also demonstrate a need for more education on the generative AI product landscape for higher education.

# Acknowledgments

We are very grateful to the participants in our cohort project, listed in Appendix A, who made this research possible.

We also express our thanks to Gary Price of Library Journal's infoDOCKET for invaluable help in keeping us up to date on the generative AI and higher education landscape.

# Methods

This study was conducted in collaboration with 19 institutions in the US and Canada that participated in Ithaka S+R's Making AI Generative for Higher Education cohort project. For a list of participating institutions, see Appendix A. Ithaka S+R developed a set of interview questions asking about experiences with generative AI in teaching and research contexts, as well as support needs. A copy of the interview guide is included in Appendix C. Interviewees were individuals with both teaching and research responsibilities at their institution, e.g., faculty, postdoctoral students, and graduate students. Each participating institution in the cohort formed a research team, received interview training from Ithaka S+R, and conducted semi-structured interviews at their respective institutions using Ithaka S+R's interview guide. Ithaka S+R also conducted supplemental interviews to enhance disciplinary diversity. The interviews were conducted between March and May 2024.

Each institution's team and Ithaka S+R conducted an average of 12 to 13 interviews, for a combined total of 246 interviews. Institutional teams were not given restrictions in their selection of interviewees in terms of department or rank; the interviews represent a wide range of disciplines and career stages. Ithaka S+R organized interviewees' ranks and disciplines into standardized categories (the standardized ranks and disciplinary affiliations of interviewees are listed in Appendix B). Two Ithaka S+R analysts used a sample of five transcripts to develop qualitative codes using a grounded theory approach, checking inter-analyst agreement. A representative sample of 45 interviews was selected for analysis. Qualitative thematic coding and analysis of the full sample in NVivo was completed by one analyst.<sup>12</sup>

---

<sup>12</sup> In addition to the typical limitations of qualitative research (e.g., findings that are directional rather than representative; interpretive bias), this study had the following limitations: 1) Automated transcriptions often contained transcription errors that could not be corrected. As a result of these limitations, precise comprehension of certain sections of these interviews was not possible. Incomprehensible sections of transcripts due to automated transcriptions were disregarded for our study. 2) Some interviewees melded their discussion of generative AI with other forms of AI, especially when it came to uses for their research. To keep this report's focus on generative AI, our analysis focused on interviewee comments about generative AI, not other forms of AI, to the best of our ability. 3) Metadata included the interviewee's department or rank, which provided



Direct quotations from interviews have been lightly edited for clarity.

# Findings

## Teaching and Learning

Whether they are optimistic or pessimistic about AI's impact, instructors across the board agreed that AI's increasing integration into teaching and learning activities feels inevitable. As an associate professor in a language department put it, "AI is not going anywhere." Interviewees also acknowledged that they cannot ignore the challenges posed by the advent of generative AI: they have to, in the words of one instructor in film studies, "avoid the head in the sand" and face the issue head-on.

This was particularly the case when it came to integrating AI into student learning. Because students are actively using AI—and will need to know how to use it for their future careers—most instructors had at least tried out using AI for student activities. Some had also experimented with using it to streamline their own teaching workflows. However, uses were largely exploratory and geared around improving student and instructor AI literacy levels: there is still significant progress to make in establishing how generative AI can be used responsibly and effectively in specific courses and disciplines in the longer term.<sup>13</sup>

In order to help users think more critically about the best uses for generative AI, it will first be necessary to continue raising individuals' levels of familiarity with the technology. Skeptics with low adoption levels also tended to have low levels of familiarity. Overall, interviewees in our

---

variable degrees of specificity about their discipline. Disciplinary distribution for our sample was determined as accurately as possible based on the available information.

<sup>13</sup> Ithaka S+R's 2024 National Instructor Survey found that 72 percent of instructors had experimented with using generative AI as an instructional tool, but only 14 percent either agreed or strongly agreed that they were confident in their ability to use generative AI in their teaching abilities. This reinforces what was reflected in our interviews: exploration with AI is widespread, but many fewer instructors have determined how to best apply the technology. See Dylan Ruediger, Melissa Blankstein, and Sage Love, "Generative AI and Postsecondary Instructional Practices: Findings from a National Survey of Instructors," *Ithaka S+R*, June 20, 2024, <https://doi.org/10.18665/sr.320892>.

study had widely varied levels of familiarity. Having a “medium” level of familiarity was most common. For certain disciplines, average familiarity levels were slightly skewed towards one end of the scale or the other. STEM interviewees, for instance, tended to have high levels of familiarity. However, an unusually large number of STEM interviewees in our study described themselves as specialists in AI, even if not in generative AI specifically, which is likely a reason behind these particularly high familiarity levels. Additionally, these interviewees reported that many of their departmental colleagues were not as familiar with AI as they were, indicating that within STEM there is also significant variance in faculty familiarity levels.

In the social sciences, interviewees fell evenly across the spectrum of familiarity levels. Individuals from the arts tended to have low levels of familiarity, though there were notable exceptions of individuals working at the intersection of technology and art who were highly engaged with AI. More humanities interviewees had medium to high levels of familiarity with generative AI than low levels. This was due, most commonly, in reaction to the discourse around student writing and generative AI. In essence, humanities instructors had a strong impetus to familiarize themselves quickly with generative AI to keep up with their students and make informed decisions on how to handle writing-based assessments. Even outside of the humanities, keeping up with students was a significant motivating factor for instructors to begin incorporating generative AI into their teaching activities.

## **Adoption in Teaching and Learning**

### **Learning Objectives and AI Literacy**

Much of the conversation around integrating generative AI into student learning came back to learning objectives. Instructors are asking themselves: what is most important for my students to learn? Can generative AI be leveraged to serve my learning objectives? Is generative AI inhibiting those objectives, or should I change them in light of this technology’s existence?

An assistant professor in engineering summed up the dilemma many are facing, explaining that after creating a “permissible” AI policy and seeing an uptick in student use, they observed that “the average quality of work

was higher than it's been in prior semesters. That is really great.” On the other hand, “there's a downside, which is that these students didn't learn some of the things they tended to learn before. So, there's this trade-off here.” The result, they explained, is that generative AI has obliged them to “take a closer look at my learning objectives” and “then analyze, what are the things that I care about [the students learning]?” This instructor recognized their students had declining skills in writing pseudocode, for example, due to AI use, but ultimately decided, “I didn't care in the context of my course if they lost that skill.” They added, though, “That only makes sense in the context of my course though, whereas in other courses, I would not be okay with that.” They emphasized that generative AI is “impacting what the students are learning” and that instructors may “need to change what [they] are doing to adapt to this.” In other words, instructors might be best served in altering their learning objectives in certain cases, asking themselves the question of what skills or knowledge are most worth teaching given AI's existence.

### The most common way in which instructors reported integrating generative AI into student coursework was through AI literacy-oriented activities.

For many instructors, familiarizing their students with AI has *become* one of their learning objectives, because of how important they think it is for students to be proficient in AI for their future careers. As an assistant professor in medicine explained, “It's not like AI is going to replace humans. It's just an expert who knows AI is going to replace an expert who doesn't.” Whether personally excited by generative AI or not, such instructors are embracing it out of a sense of duty to their students. One associate professor in education remarked, “I teach either pre-service teachers or current in-service teachers and AI is essential for them... I think part of my job as a faculty member is to teach them how to use it accurately, efficiently and ethically, so that they can then incorporate it into their own teaching practices.” Along similar lines, a political science professor stated: “I think we would be failing our students if we don't actually provide them with the critical skills that they're going to need to use the tool well.” Often, the focus on student career preparation came from instructors teaching their students content geared towards preparing them for a specific profession, such as in the cases of the instructors from schools of medicine and education cited above.

As a result of this interest in preparing students for an AI-infused world, the most common way in which instructors reported integrating these tools into student coursework was through AI literacy-oriented activities, which is to say, activities in which the instructor has students use generative AI with the objective of increasing students' familiarity with and understanding of the tools' capabilities and limitations. Recent survey data suggests that this still remains a highly popular AI use for teachers.<sup>14</sup> AI-literacy-oriented activities usually involve students creating content with generative AI tools, then critically evaluating that content under instructor supervision. For instance, a writing instructor reported using AI to summarize course content, then had their students evaluate the quality of the summaries. A professor in health sciences had students generate research reports then identify hallucinated citations. An instructor in the social sciences had AI produce code to solve a problem the students had already written code for on their own, then had the students compare their work to the AI output, analyzing why the training data led the tool to provide "inefficient solutions."

Instructors who had tried out these AI-literacy oriented activities viewed them as highly successful. Instructors across disciplines expressed doubts that their students have the literacy levels to interact safely and productively with generative AI on their own; performing these exercises in class mitigated some of those instructor worries. Not only did they feel their students' understanding of the strengths and pitfalls of the tools was improving, but instructors themselves reported learning from these experimental activities alongside their students. Whether activities primarily geared towards building up critical AI skills will have long-term value is hard to say. One professor of political science expressed the concern that such activities might not be useful forever, remarking:

some people were very excited... because they could make an assignment, and ChatGPT would come up with all the wrong information, and so, then, students would see how ineffective a tool it was, or something like that. I mean, you can do that once, so

---

<sup>14</sup> The Digital Education Council's 2025 Global AI Faculty Survey found that 50 percent of faculty respondents were "teaching students to use and evaluate AI in class." The only more common uses of AI in teaching and learning were creating teaching materials and support for administrative tasks. See "Digital Education Council Global AI Faculty Survey 2025," *Digital Education Council*, January 20, 2025, <https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-faculty-survey>.

that's not like any kind of ongoing use of those tools. And the tools are going to get better, and so, that's not going to be the case over time.

Indeed, many instructors presented their AI literacy-oriented activities as one-off experiments to familiarize themselves and their students with the new tools at their disposal. Providing opportunities for both students and faculty to critically assess the technology's potential, limitations, and risks will continue to be an important learning objective. But as common user knowledge of AI improves and new AI tools with different capabilities are released, AI literacy will be a moving target, requiring continual adaptation by instructors.

Finally, some instructors also described the opportunity to demonstrate just how important the skills they teach are, particularly in the age of AI. This was seen most among humanities interviewees, who argued that the critical reading skills they aim to teach are essential for students to evaluate the accuracy and biases of AI outputs. One history instructor, for instance, compared “what historians do with historical sources” in terms of “dissecting” and “cross referencing” to the similar skills students need to evaluate AI-generated text. Writing instructors reported seeing generative AI as an opportunity to teach students to distinguish different kinds of writing and understand that their objective in teaching writing is to teach students to learn and think through writing. One English instructor who teaches first-year writing, for instance, described using AI-generated writing to point out the difference to their students between formulaic writing that comes up with simple answers, and writing as a tool to think in critical and complex ways about a subject. Such instructors expressed hope that this kind of approach may be able to mitigate the widely touted threat generative AI poses to student critical thinking skills by making students apply those very skills when working with generative AI.

## **Student Coursework**

Our study did reveal several notable examples of instructors who had begun to identify ways to leverage AI to meet their course's core learning objectives. For instance, a social science instructor had their students evaluate AI-generated code to facilitate their process of learning programming. This was particularly useful, they argued, when teaching students without a computing or mathematics background. As they explained:

I found that using some of these generative models like Copilot, GitHub Copilot, kind of help students learn programming without getting bogged down in syntax errors. So when you're a beginner programmer, you end up spending like 20 hours for a problem that would take an expert programmer 20 minutes just because you're caught up in all these tiny little syntax errors ...So by interacting with a lot of these generative models, many teachers other than myself as well have found that it helps students kind of better situate themselves with the programming language.

After having students compare the code they wrote with AI-generated code, the instructor had students annotate the AI-generated code, noting where and how it could be improved, as a way of helping them learn better coding themselves. Such activities have the potential to become longer term teaching practices, particularly if they continue to allow instructors to get students to meet learning objectives more easily—in this case, teaching coding.<sup>15</sup>

Other instructors found it valuable to encourage students to use AI within research workflows. For instance, a professor in health sciences had students produce AI summaries of 60 articles to facilitate the process of choosing two to three articles that interest them to work on for an assignment. A psychology instructor had students working on group projects use AI to brainstorm project ideas. Where relevant, instructors also tried incorporating generative AI into practical training for students' future careers. Examples included having students in a marketing course brainstorm with AI on how to handle a social media crisis, having business students use AI to generate an image of a product for a presentation pitching it to investors, and having law students use it to help prepare for a negotiation in family law.

While wider discussions around AI pedagogy have often mentioned its potential for personalized learning experiences, such as using chatbots as “personal tutors” for students, relatively few interviewees in our study reported creating specific assignments encouraging students to leverage

---

<sup>15</sup> For a recent study of AI's potential in coding education, see Allen Nie, Yash Chandak, Miroslav Suzara, Malika Ali, Juliette Woodrow, Matt Peng, Mehran Sahami, Emma Brunskill, and Chris Piech, “The GPT Surprise: Offering Large Language Model Chat in a Massive Coding Class Reduced Engagement but Increased Adopters' Exam Performances,” OSF Preprint, April 24, 2024, last edited May 26, 2024, <https://doi.org/10.31219/osf.io/qy8zd>.

AI in this way.<sup>16</sup> Activities oriented towards building up students' critical AI skills were the much more popular category of use case. This demonstrated how AI literacy is the first priority for instructors—many expressed doubts over their students' abilities to critically evaluate outputs, understand privacy risks, or recognize the line between using AI as an assistant and “over-using” it. Once instructors feel that their students have the skillset to use AI responsibly, they will be more likely to facilitate opportunities for students to make use of AI for personalized learning.<sup>17</sup> Faculty trust in allowing students to use AI as a personalized tutor will also increase in instances where institutions can offer their faculty opportunities to use vetted generative AI tools trained specifically on course content.<sup>18</sup> In either case, advancing students' AI literacy levels is a crucial endeavor towards enabling further integration of generative AI into student coursework.

## Instructor Workflows

Incorporating generative AI into student activities was the priority for most instructors: action in that space could not wait, as instructors recognized their students were increasingly using the tools and needed the critical

---

<sup>16</sup> For discussions of generative AI's potential for personalized learning or as a personal tutor, see, for example, Xibing Wang, Xiaoshu Xu, Yunfeng Zhang, Shanshan Hao, and Weng Jie, “Exploring the Impact of Artificial Intelligence Application in Personalized Learning Environments: Thematic Analysis of Undergraduates' Perceptions in China,” *Humanities and Social Sciences Communications* 11, no. 1 (2024): Article 1644, <https://doi.org/10.1057/s41599-024-04168-x>; Megan Morrone, “AI Tutors Are Already Changing Higher Ed,” *Axios*, October 29, 2024, <https://www.axios.com/2024/10/29/ai-tutors-college-students-efficiency>; Patrick Boyle, “AI in Medical Education: 5 Ways Schools Are Employing New Tools,” *AAMCNews*, February 27, 2025, <https://www.aamc.org/news/ai-medical-education-5-ways-schools-are-employing-new-tools>; “Wiley & Fulton Schools of Engineering at ASU Collaborate to Develop AI Tutor,” *Wiley Newsroom*, October 24, 2024, <https://johnwiley2020news.q4web.com/press-releases/press-release-details/2024/Wiley-Fulton-Schools-of-Engineering-at-ASU-Collaborate-to-Develop-AI-Tutor/default.aspx>.

<sup>17</sup> For a recent discussion of why AI literacy is essential for students to use the technology effectively, as a personal tutor or otherwise, see Beth McMurtrie, “Should College Graduates Be AI Literate?” *The Chronicle of Higher Education*, April 3, 2025, <https://www.chronicle.com/article/should-college-graduates-be-ai-literate..>

<sup>18</sup> See, for example, the collaboration between Google and the Stephen M. Ross School of Business at the University of Michigan to launch a Virtual Teaching Assistant pilot program, leveraging agentic AI: “Google Public Sector Helps Enhance Learning at the University of Michigan with Pioneering New Agentic AI Virtual Teaching Assistant,” *Michigan Ross News*, April 7, 2025, <https://michiganross.umich.edu/news/google-public-sector-helps-enhance-learning-university-michigan-pioneering-new-agentic-ai>.



skills to do so effectively. However, instructors from all disciplines also reported experimenting with generative AI to assist them in their own workflows, particularly when it came to course preparation or student feedback.

Interviewees most commonly shared that they leverage AI to help create activities or assignments, whether to design an in-class group activity or draft a take-home practice problem set. Recent survey data confirms the continued popularity of these use cases for instructors.<sup>19</sup> One writing instructor, for instance, stated that using AI to create student activities has “spiced up my classroom.” Interviewees also described using AI for lecture preparation, particularly to summarize information or create bullet points. Rendering complex ideas in an easily digestible format for students was usually the priority. One assistant professor in health sciences, for example, described using AI to come up with different analogies to explain concepts for their biology students versus their computer science students.

One writing instructor, for instance, stated that using AI to create student activities has “spiced up my classroom.”

When it came to student assessment, instructors have created rubrics with AI assistance, as well as used AI to revise their feedback on student work. A cinema instructor, for example, reported using AI to take their “messy notes” and turn them into “a more formatted version [of feedback]... or use it to build out questions that I would follow up with [the students].” On the other hand, having AI tools generate automated feedback for students was not a commonly discussed use of the technology among interviewees—their primary interest was having AI

---

<sup>19</sup> Seventy-five percent of faculty respondents to the Digital Education Council’s 2025 Global AI Faculty Survey reported using AI to create teaching materials, the most popular use of AI in the teaching and learning context; see “Digital Education Council Global AI Faculty Survey 2025,” *Digital Education Council*, January 20, 2025, <https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-faculty-survey>. It was also the most popular use case for instructors in Ithaka S+R’s 2024 National Instructor Survey, at 22 percent; see Dylan Ruediger, Melissa Blankstein, and Sage Love, “Generative AI and Postsecondary Instructional Practices: Findings from a National Survey of Instructors,” *Ithaka S+R*, June 20, 2024, <https://doi.org/10.18665/sr.320892>.



revise feedback they had written themselves.

Instructors stated their commitment to ensuring the AI-generated material met the standards to which they had previously held themselves for their courses. AI was frequently described as an “assistant” or “collaborator” when it came to teaching tasks, but one that needed significant supervision. One writing instructor echoed the feelings of many interviewees: “it definitely takes a lot of very specific prompting to kind of get what you want...but typically I'll do that and then take that idea and modify it for what works best for my classroom.” Many instructors found AI tools would only give them a base structure for course material that they would build on. Nonetheless, from the perspective of these adopters, just having AI create that base structure was already a significant time-saver.

## Non-Adoption

While most interviewees had at least dabbled in generative AI in a teaching and learning context, some had not yet applied the technology to teaching and learning at all. The most common reason for this was the learning curve: these instructors had not yet made or found the time to familiarize themselves with generative AI to the degree necessary to implement it into their work. In many cases, these instructors had not been teaching any or many classes since generative AI's emergence, due to sabbatical leave or holding leadership positions in university administration. As a result, they felt less pressure to upskill quickly than their peers who were teaching regularly. However, the vast majority of these non-adopters recognized the need to familiarize themselves with AI soon and planned to integrate AI into their practice in the near future.

The vast majority of these non-adopters recognized the need to familiarize themselves with AI soon and planned to integrate AI into their practice in the near future.

A common predicament for non-adopters—as well as adopters—was that they were open to further integrating AI but felt they had not yet figured out how or if they could productively do so in a way that would support their core learning objectives. As an assistant professor in business remarked, “I teach some coding classes. So, for me, I think that's still such

a fundamental concept. ... I don't want my students to overly use [generative AI] to circumvent their learning. I want them to use it to support their learning. But I think that's the thing that we're all struggling to figure out how to do." This fear of students no longer meeting important learning objectives as their AI use increased was a major barrier for instructors that had not identified longer term use cases for students beyond AI literacy-oriented activities. For example, within the field of law, a professor expressed concerns about how generative AI shortcuts for certain aspects of law students' training might have detrimental effects on the field:

My law students read on average probably about, in my class they'll read 60 to 100 pages a week.... This is my case book that I used for family law. It has 1,100, 1,160 pages.... They got to know the book. What happens when we have a legal community that hasn't put all those words in their head? I think about that. When I talk to attorneys who've been practicing 40 years, they can look at a contract in the blink of an eye, say that clause won't work. Why? Because they've looked at thousands and thousands and thousands of contracts over their practice. They can see in the blink of an eye a harm, a failure, the implications because they've learned it by what works, what doesn't work... What happens [when] we have a generation that's [having] AI tell it what would be a good contract and we run into confirmation bias.... I worry and I raise that to judges. I raise it to the legal community.

Such comments reveal the uncertainty many still feel about the longer-term impacts of AI-infused student learning. As a result, making decisions about how and if they can productively integrate generative AI into their courses still feels daunting for many instructors.

More common than flat-out non-adopters were instructors who had integrated generative AI into certain aspects of their practice but not others. In some cases, they planned to expand their use of AI in the future. In other cases, they saw AI as inherently limited to assisting with specific teaching or learning tasks. A math instructor, for instance, allows their students to use AI, and describes its impact on teaching and learning as "transformative." However, when it came to using generative AI to create materials for teaching, they still found it more efficient to do that work themselves. As they explained, "I have a lot of lecture notes that are historically archived... And those old lectures have through years of

polishing, and correction, and error checking, and so forth, so I think that the quality of most of those teach[ing] materials is better than I can currently get from GPT-4.”

## Academic Integrity: Concerns and Policies

In the immediate wake of ChatGPT’s commercial release, the issue of students breaching academic integrity policies was top of mind across the higher education sector. Since then, many instructors have been working through more nuanced questions of how to approach academic integrity with their students and craft effective policies.

Interviewees were firm in believing that generative AI could never be effectively policed. They found it difficult, if not impossible, to prove a student had used a generative AI tool and suspected this would only become more challenging as the technology continues to improve. Instead, the majority expressed a desire to build up a mutual understanding with their students around appropriate and inappropriate uses of generative AI so that the students would, in effect, “police” themselves. Many expressed wariness about the adversarial relationship they would create with students if they became obsessive about detecting AI use. A theater instructor explained, “I definitely don’t want to have a classroom environment where AI has created the sense that I can’t trust you, that I’m constantly monitoring you, that I’m always questioning your work. I want to operate on the assumption that people are generally trustworthy, and they’re doing their best.” When catching a student breaching policies, some instructors wanted to take it as a chance to, as one English instructor put it, “engage them directly in talking about their process and turn it into a learning opportunity for the student.”

“I want to operate on the assumption that people are generally trustworthy, and they’re doing their best.”

However, even instructors who insisted on their belief in fostering mutual understanding around generative AI with their students sometimes found it challenging to implement that in practice. For instance, a writing instructor with a policy permitting certain uses and requiring documentation caught students breaching it. They created a more

detailed policy, then caught more students. The frustrated instructor then banned AI altogether in the course. As they explained:

I got increasingly mad when it happened multiple times. And by the time four people had done it in the fall semester, like a month in, I was so annoyed. I said, “Okay. So, if some people have ruined it for everyone else, then now there's a ban unless you have permission.” ... And I was annoyed with them and also annoyed with myself because I don't necessarily believe in an outright ban. But to me, it was exhibiting the sense that I think they were aware they were over relying on it, which is why they weren't disclosing it because none of them were disclosing it. And that's what I was most annoyed about. I'm like, “Why aren't you documenting it? I told you how to do that.”

The instructor explained they have not banned AI in other courses but instead have tried to have “more detailed conversations about AI” in class and have created activities where students critically evaluate AI outputs under instructor supervision. In sum, even if the panic around student academic integrity has subsided from its previous levels, many instructors still feel they have yet to establish the best approach with students—and it may be a long process of trial and error to get there.<sup>20</sup>

Instructors who had a course policy usually followed the formula of permitting certain AI uses but required students to document usage. Policies often specified that AI could only be used as a “collaborator” or “assistant.” They also tended to state that only certain use cases were acceptable—such as ideating or revision—but students could not copy and paste generated text into assessed work. A few instructors had created more innovative versions of this type of policy, such as one from an English department who maintains a regularly updated “Frequently Asked Questions section” of a course syllabus, which contains student questions about appropriate AI use cases and instructor responses.

---

<sup>20</sup> The American Association of Colleges and Universities’ (AAC&U) recent survey of senior leaders at higher education institutions found that the majority (59 percent) of respondents have seen an increase in cheating since generative AI tools became widely available, indicating that faculty are still struggling to enforce academic integrity policies in generative AI’s wake. See C. Edward Watson and Lee Rainie, “Leading Through Disruption: Higher Education Executives Assess AI’s Impacts on Teaching and Learning,” *American Association of Colleges and Universities*, 2025, <https://www.aacu.org/research/leading-through-disruption>.

Instructors who did not have a formal policy in their syllabus were usually handling AI use in the classroom on a case-by-case basis, either when they suspected students might use it for a specific assignment or when approached by a student asking for permission. Many instructors were relying on in-person exams, but extremely few reported outright bans. Most instructors without a policy recognized they would need to create one soon.

## **Aligning Policies and Strategies**

The vast majority of instructors reported that their university did not have a prescriptive stance on generative AI use in the classroom. Their institution provided them with sample options of syllabus language but left instructors full autonomy in choosing how and if to integrate generative AI into their courses. This remains the dominant approach from university administrations. While many interviewees expressed their appreciation for this open stance, others pointed out its pitfalls—namely, the inconsistency in policies and adoption levels across instructors and disciplines, resulting in confusion and uneven levels of exposure to the technology among students.

While instructors often pointed to the need for institutions to have a more unified approach to AI's role in teaching and learning, they also felt that policies and adoption would have to vary by discipline, making this a sticky issue to resolve. For instance, an instructor from a film department suspected, like many others, that their students were confused by policy variance: “they may find something [that they can do with AI] they like in one class that they want to apply to another class, and they don't know if they should, or they don't know if they can, or they don't think it's wrong.” Part of the issue, one engineering assistant professor pointed out, is that many instructors do not even have a clear policy. As a result, some students are using AI with full freedom, and others are afraid of using it in a way that would breach academic integrity standards—standards that are not clear to them. As this instructor explained, “I do think we need to standardize what's allowed in a class and make it so there's almost a few standard AI policies and we follow one of them, to allow students to understand this more consistently.”<sup>21</sup>

---

<sup>21</sup> Inside Higher Ed's 2024 Student Voice survey found that three in 10 students are not clear on when they are or are not allowed to use generative AI in their coursework, indicating a need for further guidance; see Ashley Mowreader, “Survey: When Should College Students Use AI? They're Not Sure,” *Inside Higher Ed*, September 16, 2024,

Many interviewees advocated for a more formalized, either departmental or institution-wide approach to integrating AI into student learning, especially basic AI literacy skills. One writing instructor echoed the comments of many: “if we’re talking next on the agenda, it’s really how we develop that AI literacy, almost like any other skill we develop, and that’s part of the general education.” However, there was little agreement on exactly how this would be done and whose responsibility it would be—that of individual instructors, departments, specific units, or senior leadership-led initiatives. Suggestions ranged from making AI literacy part of first year writing courses, or a part of methods courses in each department to allow for discipline-specific nuances to be addressed. Others thought the topic should be covered in a required training module as a part of new student onboarding. Recent data suggests that there remains significant room for growth when it comes to integrating AI literacy into general education curricula.<sup>22</sup>

Our findings underscore several important issues for the immediate future of teaching and learning at universities. How can instructors integrate generative AI into student learning while still meeting learning objectives—and what should those objectives even look like in the age of generative AI? How can institutions foster coherence across course policies, while also making room for disciplinary differences and respecting the autonomy of instructors? What can institutions do to build a bridge between individual assignments and a systematic integration of AI literacy into students’ educational experiences? These are questions that our

---

<https://www.insidehighered.com/news/student-success/academic-life/2024/09/16/college-students-uncertain-about-ai-policies>. See also Inside Higher Ed’s discussion of the “wild west” of varied or nonsubstantive AI policies at universities: Kathryn Palmer, “Is Grammarly AI? Notre Dame Says Yes,” *Inside Higher Ed*, November 26, 2024, <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2024/11/26/grammarly-ai-notre-dame-says-yes>.

<sup>22</sup> Fifty-eight percent of student respondents to the Digital Education Council’s 2024 Global AI Student Survey felt they do not have sufficient AI knowledge and skills, and 72 percent thought that universities should provide training for students on the effective use of AI tools; see “Digital Education Council Global AI Student Survey 2024,” *Digital Education Council*, August 2, 2024, <https://www.digitaleducationcouncil.com/post/digital-education-council-global-ai-student-survey-2024>. However, only 14 percent of respondents to AAC&U’s recent survey of senior leaders at higher education institutions said they had set AI literacy as a general education learning outcome at their institution; see C. Edward Watson and Lee Rainie, “Leading Through Disruption: Higher Education Executives Assess AI’s Impacts on Teaching and Learning,” *American Association of Colleges and Universities*, 2025, <https://www.aacu.org/research/leading-through-disruption>.

interviewees were well aware of—and in some cases, had even made headway on—but to which they did not yet have firm answers.

## Research

Many instructors felt obligated to immediately delve into generative AI in the teaching and learning context to keep up with their students. The urgency was not the same for incorporating generative AI in their research work. Nonetheless, over half of interviewees reported having at least tried out generative AI for research-related tasks. Generally speaking, interviewee comments on AI in research contexts were not as nuanced as those on teaching and learning, showing that the depth and urgency in conversations about AI in research still lag slightly behind those about AI in teaching and learning.

**Use of generative AI in research contexts was often minimal and experimental.**

As with teaching—and perhaps even more so in the case of research—interviewee use of generative AI in research contexts was often minimal and experimental.<sup>23</sup> While some individuals were ahead of the curve and described themselves as heavy users, most had not yet identified the best practices and uses of generative AI in their research workflows for the longer term. Clear trends emerged from our study in terms of which stages of research generative AI was being tried out or applied: interviewees were most likely to either be using generative AI in early stages of their process—such as for brainstorming or outlining—or at the very end, especially for revising writing. Interviewees also reported using generative AI tools to summarize scholarship and assist with literature, though there were varied opinions on how useful generative AI was in those specific areas.

Interviewees tended to have an optimistic outlook on generative AI's potential to accelerate and improve research. At the same time, they

---

<sup>23</sup> Ithaka S+R's 2024 biomedical researcher survey found that 63 percent of biomedical researchers have used generative AI in their research, but the vast majority are not using it regularly. The numbers for researchers surveyed outside biomedicine were similar. See Dylan Ruediger, Chelsea McCracken, and Makala Skinner, "Adoption of Generative AI by Academic Biomedical Researchers," *Ithaka S+R*, October 17, 2024, <https://doi.org/10.18665/sr.321415>.



showed strong awareness of AI's limitations and expressed concerns about upholding research quality and integrity. When researchers adopted generative AI, it was because they felt reasonably confident they were not breaching ethical standards and it had proved to be helpful for what they saw as challenging, onerous, or mindless tasks. However, this has left the research community with a series of challenging questions, namely: Which aspects of research need to be done by a human for the research to count as that individual's work? How can generative AI be effectively leveraged in research contexts without significant risk to the quality and integrity of the research product, and how and when will discipline-wide standards be set for this? Like in the teaching and learning context, in order for generative AI to be productively and responsibly mobilized in research contexts, the research community must make a successful transition from experimenting with the technology to fostering strong AI literacy skills in researchers, especially when it comes to thinking critically about the technology to identify ethics and best practices. Researchers interviewed for our study also expressed a strong desire for more guidance on ethics and best practices from within their fields and from the academic community at large.<sup>24</sup>

## **Adoption in Research**

### **Brainstorming, Outlining, and Discovery**

Brainstorming and organizing ideas, when beginning a research project or a piece of writing, were particularly popular contexts in which to apply AI. This was especially the case for researchers in the humanities and social sciences. These interviewees found brainstorming and idea organization important yet time-consuming parts of their work, where generative AI could be brought in as an assistant without feeling like it was doing too much work for them. For one interviewee from an education department, for instance, ChatGPT had become like “a member of the brainstorming team.” As they described it: “I mainly used it as a collaborator in that if I

---

<sup>24</sup> Ithaka S+R's 2024 biomedical researcher survey found that the most common barriers for biomedical researchers to incorporating generative AI into their research were insufficient levels of accuracy and/or reliability in generative AI outputs, and lack of clarity about best practices for research integrity while using generative AI; see Dylan Ruediger, Chelsea McCracken, and Makala Skinner, “Adoption of Generative AI by Academic Biomedical Researchers,” *Ithaka S+R*, October 17, 2024, <https://doi.org/10.18665/sr.321415>. Similarly, Wiley's 2024 ExplanAltions study found that 63 percent of researchers reported their use of AI was inhibited due to a lack of guidelines or training; see “ExplanAltions: An AI Study by Wiley,” *Wiley*, February 4, 2025, <https://www.wiley.com/en-us/ai-study/for-researchers>.



were to say, alright, we need to come up with this research question. What are you thinking? This is what I'm thinking, and now let's go ahead and take this all and edit it until we get what we want.” Along similar lines, an instructor in cinema studies described using ChatGPT’s voice capabilities to chat while walking the dog in the morning, having “very fluid conversations just to work out an idea.” Having generative AI produce outlines was also a popular use. One law professor, for instance, described providing ChatGPT with the information they want to present and the target audience, then asking for suggestions on what to prioritize and how to format the presentation.

Interviewees who reported success using generative AI to brainstorm or outline did so while critically engaging with its output. Often they were not entirely satisfied with what the AI tool suggested. Nonetheless, as one researcher in education put it, “it gets the wheels turning...it's a great starting point.” Much like in the teaching context, the interviewees who were most satisfied with their experiences using AI in their work were those who were willing to take the time to carefully craft and revise prompts, and those who built from AI outputs as starting points rather than using those outputs directly.

Beyond ideating and organizing ideas, generative AI was also used in the early stages of the research process when it came to search and discovery. Researchers described trying it out both for locating scholarly sources of interest, as well as for general information searches. However, opinions varied widely on the usefulness of AI in these contexts. It was not the most popular use case for research in any field—with hallucinations being the main issue—though STEM researchers were more likely to report using AI in these ways than other disciplines. Those who did use generalist tools to search for sources or information were aware of the risk of inaccuracies and trusted themselves to judge the quality of the outputs. Most interviewees, however, admitted they still relied primarily on Google, Google Scholar, and their library catalog for searching, even if they had tried out generative AI for similar purposes. As will be further discussed later in this report, the majority of interviewees were using ChatGPT, rather than tools grounded in vetted scholarly content, meaning that the potential utility of generative AI in the discovery phase of research was widely untapped.<sup>25</sup>

---

<sup>25</sup> For a discussion of the unreliability of LLMs like ChatGPT in academic writing in comparison to scholarly databases, see Swati Garg, Asad Ahmad, and Dag Øivind

## Data Collection and Analysis

Our study revealed budding generative AI use to assist with collection or analysis of qualitative or quantitative data, even if such use cases were not as frequently mentioned by interviewees as brainstorming, summarizing, or revising writing. This category of use cases was more often reported by interviewees from STEM and social science fields than in those from the arts and humanities. One doctoral student using social science methodologies in the field of communication, for example, described how they had leveraged AI tools when it came to conducting interviews:

So, in my research I've done interviews and so I've found AI very useful in transcribing the interviews, creating summaries of interviews... It's incredibly helpful with coming up with interview questions. So, if I'm going to do an interview, I'll interact with the AI ahead of time in order to make myself a little bit more prepared... The times that it hasn't been helpful are the times when I've tried to get it to make some decisions and do my work for me... it was okay at getting some themes out of the transcripts, but it wasn't good at highlighting good quotes. It wasn't good at selecting specific pieces of the transcription that I would use for my work. I still had to do that.

This example demonstrates how generative AI has the potential to assist in the creation of data collection instruments, even if just as a starting point to iterate on with clear limitations. This description highlights a common experience among researchers: finding that generative AI helps with some, but not all, aspects of their work. The key was discovering its strengths and acknowledging its weaknesses. For this interviewee, the tool's strengths were its ability to help brainstorm interview questions, but it yielded poorer results when it came to deciding which quotes from the interviews were most meaningful. This PhD student was not alone in finding that AI was weak in making judgment calls about what information

---

Madsen, "Academic Writing in the Age of AI: Comparing the Reliability of ChatGPT and Bard with Scopus and Web of Science," *Journal of Innovation & Knowledge* 9, no. 4 (2024): Article 100563, <https://doi.org/10.1016/j.jik.2024.100563>. For a discussion of the breadth of discovery-focused AI tools for the higher education market, see Claire Baytas and Dylan Ruediger, "Generative AI in Higher Education: The Product Landscape" *Ithaka S+R*, March 7, 2024, <https://doi.org/10.18665/sr.320394>.

was meaningful or significant.<sup>26</sup>

Additionally, interviewees reported using generative AI to generate, debug, or ask questions about code. This was one of the most popular categories of use among STEM interviewees, and interviewees in other disciplines used generative AI for this purpose as well. In addition to speeding up workflows, generative AI's proficiency in coding also lowers the barrier for less experienced coders. By easing the learning curve, generative AI may open doors for certain researchers to integrate coding into their research methodologies who would not have done so previously.

## **Summarization and Literature Reviews**

Researchers were experimenting with generative AI's ability to summarize information in the context of summarizing existing research for literature reviews, generating abstracts, and learning about a subject area outside of their area of expertise. However, there were vastly contrasting opinions on the utility of using AI-generated summaries in these ways due to concerns about inaccuracies. Researchers in STEM were more likely to have reported using generative AI for summarization and literature review assistance, but there were adopters and non-adopters across disciplines.

Opinions around the ethics and usefulness of generative AI in these contexts circled back to a few key questions. Does using generative AI to summarize existing scholarship mean scholars are not reading and understanding that scholarship in the same way as before? Does the act of summarizing others' work involve creativity and critical thinking, or is it the equivalent of busy work? If the former, does that mean generative AI should not be used to summarize others' work? The variety of opinions our study revealed when it came to using generative AI for literature reviews demonstrates that these are questions to which the research community still has mixed responses.

Researchers who were skeptical about using generative AI to help write literature reviews were concerned about the potential negative effect on research quality. One associate professor in education, for instance, recounted an anecdote that several other interviewees also shared: they had AI generate a summary of their own work, were disappointed in the

---

<sup>26</sup> On the limitations of using AI in qualitative research, see Andrew L. Gillen, "Can We Trust AI in Qualitative Research?" *Inside Higher Ed*, October 9, 2024, <https://www.insidehighered.com/opinion/views/2024/10/09/can-we-trust-ai-qualitative-research-opinion>.

output, and as a result lost trust in the accuracy and nuance of AI-generated summaries generally speaking. They tried it out on other articles, with similarly discouraging results. As they explained, “I still had to go back to all those articles and read them, and so then it just ended up being a waste of time.” After these experiences, this interviewee explained, “for a moment, I felt a little old school, but in thinking that if I’m going to really write this full literature review and synthesize literature, I need to really read these articles and understand more than just what [generative AI tools] think are the topic sentences for all these paragraphs.” They subsequently turned down an offer from their department chair to fund a subscription to the summarization tool they had tried out.

On the other hand, for those leveraging generative AI for literature reviews, mitigating information overload was one of the main appeals. For example, one professor in engineering saw using generative AI to help keep up with the ever-increasing quantity of existing scholarly literature in today’s world as “a must.” They also offered an anecdote of a time when their research output suffered because they were unable to keep up with the existing literature. “When I was a graduate student,” they explained, “I spent maybe half a year working on a problem and I found a solution that made me so happy. A week later, when I was preparing a manuscript for publication, I found out that someone else did it before me. It was such a deflation.” In this researcher’s opinion, generative AI makes it so that “today, this is less likely to happen. It is true, we have many more papers that are being published, but on the other hand, we have these tools that can basically scan through papers and find similarities and then alert you.”

For other researchers, it was less about ensuring their research was still original, and more about using generative AI to partially automate what they saw as a non-intellectual, unenjoyable aspect of the research process. For instance, one associate professor in history, who has not yet applied generative AI to their own research, expressed their excitement about trying it out to assist with writing literature reviews:

By far my least favorite part of writing is summarizing what other people have said. I just find it completely void of any creativity. You gotta know it. So, I like learning it. But then, having to type it out for somebody else, I’m like, come on, this is not fun. So, I could really see using a tool to kind of get me over that first step to summarize

things and then go in and edit and add and delete and reframe and emphasize different things, because it's just a lot of grunt work for me.

This statement reflects common sentiments among the interviewees using generative AI for literature reviews: they maintained that knowing literature in their fields was still important but hoped to leverage generative AI to do the “grunt work”—which is to say, crafting the first draft summary of that literature.

This researcher’s comments also foreground the important question of where to draw the line between what is an unintellectual and uncreative task and what is not, and whether this varies by discipline. Researchers gain something intellectually by crafting their own summaries and contribute something valuable to the information ecosystem by writing a summary of existing knowledge shaped by their individual perspective and training. Such questions will have to be addressed to best articulate ethical standards and best practices for how and when generative AI should be used in the literature review process.

Beyond leveraging AI to conduct literature reviews within their field, a few interviewees described how generative AI could help them learn about subject matters outside of their area of expertise, thus facilitating interdisciplinary research. As a professor from a business school explained, the highly specialized nature of research makes it so that “you’re almost wearing blinders.” Generative AI, they thought, could help them think about how to “marry” two topics when they do not know much about one of them. However, other interviewees cast doubts over whether researchers could evaluate the accuracy of generative AI outputs on subjects they do not know intimately. A professor in health sciences, for instance, described their initial excitement about using generative AI to “[mesh] worlds that hadn’t meshed before,” bringing together different topics and critical frameworks. However, they began detecting falsehoods in the generative AI outputs and underlined that the next generation of researchers would need to be trained with “cautionary advice” to always double-check such outputs. While researchers may easily see inaccuracies in AI-generated summaries about their own work or field of expertise, the concern is whether they will be able to effectively vet AI-generated information every time in other contexts.

The time saving advantages of generative AI for summarization or literature reviews means that it is likely that this type of use will only become more widespread.

In an age where researchers have unprecedented access to the ever-growing quantity of existing research in their field, and the pressure to publish at scale only increases, the time saving advantages of generative AI for summarization or literature reviews means that it is likely that this type of use will only become more widespread. It is also worth mentioning that publishers, aggregators, content providers, and other vendors are making significant investments in creating tools that facilitate the process of comprehending scholarly content, with features that summarize, help automate literature reviews, or allow researchers to query documents.<sup>27</sup> In other words, current signs indicate that these use cases will likely continue to rise in popularity.

## Revising Research Outputs

Other than early stage uses such as brainstorming or outlining, researchers most commonly reported using generative AI towards the end of their research process to revise written outputs.<sup>28</sup> For instance, a researcher in earth sciences explained how they use generative AI to revise an article draft:

I'll write material, and then usually I'll only ever feed it a paragraph at a time or a section at a time, and I'll ask it to make it more readable. You'll prompt it and tell it who the audience is. So, you'd say, "this is for an audience of farmers or agricultural researchers" and then "make this paragraph more readable at a [grade] level."...

---

<sup>27</sup> See the discussion of "understanding"-oriented generative AI tools in Claire Baytas and Dylan Ruediger, "Generative AI in Higher Education: The Product Landscape" *Ithaka S+R*, March 7, 2024, <https://doi.org/10.18665/sr.320394>. See also Tracy Bergstrom and Dylan Ruediger, "A Third Transformation?: Generative AI and Scholarly Publishing," *Ithaka S+R*, October 30, 2024, <https://doi.org/10.18665/sr.321519>.

<sup>28</sup> When asked how they have used generative AI in their biomedical research, "reviewing/editing grammar" was the most popular use case (31 percent) for respondents to Ithaka S+R's 2024 survey of biomedical researchers; see Dylan Ruediger, Chelsea McCracken, and Makala Skinner, "Adoption of Generative AI by Academic Biomedical Researchers," *Ithaka S+R*, October 17, 2024, <https://doi.org/10.18665/sr.321415>.

In my experience when I do that, my paragraph and sentence structure is almost completely retained. Usually, it's about 90 percent unchanged. But that 10 percent change really helps make it more readable.

Like with other use cases, researchers understood the importance in overseeing and criticizing AI's suggestions. As this same researcher went on to explain: "what I do is I look, I'll reread what it fixed for me, and then I'll go back to my text, and I'll edit it to be more close to what it wants. And its suggestions are not always what you want, right? It always tries to generalize things, and you lose some of that specificity that you need in your writing."

Generative AI's capability to revise and grammar check was seen as a particularly effective tool for researchers writing in a non-native language. Our interviewees often discussed this issue in the context of academic journals published in English. One researcher in math, for instance, expressed how generative AI was a long-awaited solution for exceptional researchers with important knowledge to share, whose work might have been previously rejected due to concerns about language rather than the content. As they explained:

I am an editor for a journal, and about eight months ago, back in the summer, I received an article, and the math looked correct, but the writing was terrible. It was so bad I could not even send it out to review. And so I wrote back to the author and said, "Please, work with a native English speaker to turn this paper into reasonable English." And the next day, the paper came back to me in flawless English. He'd run it through GPT-4, and it had solved everything. And some of the very most brilliant minds in my field do not speak English as a first language, and that has been a career handicap to so many people. And I am delighted that GPT-4 is now removing or reducing that obstacle to so many of my colleagues.

This interviewee was not alone in arguing that generative AI would help level the playing field for non-native speakers of the dominant languages of research publications. For researchers like this one, generative AI is already breaking down previous linguistic barriers to knowledge circulation—a change that has the potential to both speed up and improve knowledge production in the field.



Using generative AI for revising writing was one of the least controversial use cases in both the research and teaching contexts among interviewees. Researchers were likely to see it as ethical and even beneficial, and, likewise, instructors were also likely to let their students use it for this purpose, especially for non-native speakers of the language of instruction. As interviewees pointed out, the question that will be challenging to answer is determining where the line is between generative AI revising writing and generative AI producing writing, especially as the latter is widely considered an unethical use case.

## Non-adoption

Researchers who reported that they have not adopted generative AI for specific tasks—or in some cases, not at all—were usually motivated by at least one of a few key reasons. As previously discussed, inaccuracies in AI outputs were a concern for researchers, in contexts such as searching for or synthesizing information.<sup>29</sup> Other interviewees argued that generative AI tools simply do not have the capabilities to assist with core parts of their research—at least, not yet. As one associate professor in chemistry noted, when they tried to use AI “to develop a more complex algorithm... it failed spectacularly. It’s not there yet.” Researchers that performed hands-on creative work, such as in the arts, or worked primarily in archives, also tended to find generative AI irrelevant to these aspects of their research. Additionally, some researchers were hesitant to incorporate generative AI in their research practice because they did not want to delegate their work to a machine. As one lecturer in psychology put it, “that feels like it’s outsourcing the intellectual labor to another entity.” A lecturer from an art department remarked, “I love research so much and I just have my own workflow that I just can’t imagine wanting to shortcut or streamline.”

Another limitation interviewees underscored was generative AI’s inability to make judgment calls. As one researcher from a math department explained:

Large language models tend not to have an appreciation of what is an interesting theorem and what is the dull theorem. There are lots

---

<sup>29</sup> Inaccuracy in generative AI is a major concern for researchers: 97 percent of respondents to Ithaka S+R’s 2024 survey of biomedical researchers felt that insufficient levels of accuracy and/or reliability in generative AI outputs was a barrier in incorporating generative AI into their own research; see Dylan Ruediger, Chelsea McCracken, and Makala Skinner, “Adoption of Generative AI by Academic Biomedical Researchers,” *Ithaka S+R*, October 17, 2024, <https://doi.org/10.18665/sr.321415>.



of trivial, dull theorems out there that nobody cares about, and that's okay, but real science, real statistics, real math advances by pursuing interesting, important problems, and right now, AIs don't have the aesthetic taste to determine what's an interesting theorem.

This interviewee points out a significant weakness of generative AI that could apply across fields: a crucial part of a researcher's work is understanding what knowledge is worth seeking and what problems are worth solving—a task that interviewees found AI had only a limited ability to help with.

A crucial part of a researcher's work is understanding what knowledge is worth seeking and what problems are worth solving—a task that interviewees found AI had only a limited ability to help with.

Across disciplines, the majority of interviewees claimed to be non-adopters when it came to generating writing in the context of scholarly publications. Interviewees were more flexible around generating content for other scholarly outputs, such as research presentations or lectures, but generating text to use in an academic journal or book was, as one law professor put it, “a line I will not cross.” While using AI to revise writing was widely accepted, researchers often described a fine but important line between using AI to revise versus to generate text. Where that line is was difficult to ascertain, but the core of the issue appeared to be whether the “ideas” expressed in academic writing came from the researcher or not. As one assistant professor in an area studies department described it, “the consensus so far in this field is that only use AI so far insofar as it doesn't change your ideas from your text. It has to come from you.” In a few other cases, researchers were reluctant to hand over writing to AI because it is a task they excel at and enjoy. One professor in engineering explained that “it's kind of funny. I actually enjoy writing even in my spare time. I enjoy writing and so it's a task that I don't really want to give away to AI and try to do it myself, because it's something I enjoy doing.”

An important caveat is that several interviewees indicated that they frequently saw AI-generated text submitted to journals in their field and suspected many of their colleagues were using AI to generate text. The existence of academic papers with undisclosed AI use has been tracked and documented.<sup>30</sup> Ultimately, it is difficult to ascertain just how many researchers are using AI to write papers in the manner most interviewees deem unacceptable. Nonetheless, our study shows there is significant hesitancy around this use case: there was no other use case where interviewees were as likely to state that it was inappropriate to employ generative AI.

## Establishing Standards and Best Practices

Similarly to how students were reportedly confused by the variance in their courses' policies, researchers were confused by what they perceived as vague policies around best practices, integrity standards, and requirements for disclosure of AI use. In order to best leverage generative AI's capabilities effectively and responsibly, researchers want more established guidelines within their fields and within the academic research community more broadly speaking.

When asked where they sought ethical guidance for using AI in research contexts, most interviewees were unable to point to a clear resource. Mostly, they made assumptions based on what they heard from colleagues or at conferences. One of their most important points of reference on best practices for research were journal policies, but interviewees reported feeling these were still unclear and inconsistent, both in terms of what AI uses are permitted and how to be transparent about use.<sup>31</sup> For example, a

---

<sup>30</sup> See Alex Glynn, "Suspected Undeclared Use of Artificial Intelligence in the Academic Literature: An Analysis of the Academ-AI Dataset," *arXiv* (2024), <https://doi.org/10.48550/arXiv.2411.15218>, and Glynn's repository of undeclared use of AI in academic literature: <https://www.academ-ai.info/>. See also Holly Else, "Should Researchers Use Ai to Write Papers? Group Aims for Community-driven Standards," *Science*, April 16, 2024, <https://www.science.org/content/article/should-researchers-use-ai-write-papers-group-aims-community-driven-standards>.

<sup>31</sup> Wiley's ExplanAltions study found that the majority of researchers are looking to publishers for guidance when it comes to AI use; see "ExplanAltions: An AI Study by Wiley," Wiley, February 4, 2025, <https://www.wiley.com/en-us/ai-study/publishers-role-ai>. Ithaka S+R's 2024 survey of biomedical researchers also found that more than half of respondents thought that explicit guidance from their publishers and funders on AI use would be helpful in incorporating AI into their research; see Dylan Ruediger, Chelsea

professor in health sciences said they felt “a little nervous” using AI to brainstorm for a journal article, even though they were not generating text for the article. The issue, they explained, is that because policies “vary so much from one journal to another one, you don’t want to accidentally block a pathway to publication just because you chose to use a tool to save you a little bit of time.”

Proponents of leveraging generative AI for research thought these confusing standards were inhibiting productive uses of the technology. For example, when asked about how their field was navigating the ethical implications of generative AI use, one researcher in a social sciences department found that unclear or strict journal policies were leading to widespread fear of using it. They thought widespread hesitancy among researchers could have the negative effect of preventing them from leveraging AI in ethical and productive ways. As this interviewee explained, “there has to be valid uses for using, like, image generation software to generate diagrams. So why should I need to develop my own diagram in PowerPoint or whatever, when I could use generative AI?”

In the past year alone, publishers, universities, and other stakeholders across higher education have increasingly begun proposing guidelines for appropriate uses, as well as frameworks for how to disclose use.<sup>32</sup> However, evidence suggests that the dust is nowhere near settled yet, as issues ranging from making publisher policies robust and detailed enough,

---

McCracken, and Makala Skinner, “Adoption of Generative AI by Academic Biomedical Researchers,” *Ithaka S+R*, October 17, 2024, <https://doi.org/10.18665/sr.321415>.

<sup>32</sup> For example, see: “Wiley Releases AI Guidelines for Authors,” *Wiley Newsroom*, March 13, 2025, <https://newsroom.wiley.com/press-releases/press-release-details/2025/Wiley-Releases-AI-Guidelines-for-Authors/default.aspx>; Amrita Ganguly, Aditya Johri, Areej Ali, and Nora McDonald, “Generative Artificial Intelligence for Academic Research: Evidence from Guidance Issued for Researchers by Higher Education Institutions in the United States,” *arXiv preprint*, arXiv:2503.00664, March 2025, <https://doi.org/10.48550/arXiv.2503.00664>; Kari D. Weaver, “The Artificial Intelligence Disclosure (AID) Framework,” *College & Research Libraries News*, 85, no. 10 (2024): 407, <https://doi.org/10.5860/crln.85.10.407>; David B. Resnik and Mohammed Hosseini, “Disclosing Artificial Intelligence Use in Scientific Research and Publication: When Should Disclosure Be Mandatory, Optional, or Unnecessary?” *Accountability in Research* (2024): 1-13, <https://doi.org/10.1080/08989621.2025.2481949>.

to figuring out how publishers can actually enforce transparency policies, remain.<sup>33</sup>

## Supporting Instructors and Researchers

As we have noted, instructors and researchers are all seeking further guidance on appropriate uses and best practices when it comes to incorporating generative AI into teaching and research. For units offering support, among the main challenges are creating support resources that cater to the highly variable levels of AI familiarity and literacy among faculty, as well as determining which resources are applicable to everyone, versus where support needs to be developed on disciplinary levels. Interviewees showed strong interest in support that applies more closely to their discipline, as well as opportunities to learn with and from their peers. They also revealed a significant need for more information about the product landscape as well as for secure and affordable access to AI tools.

## Support Resources Used

Universities are attempting to provide generative AI-related support to their faculty, often in the form of workshops, presentations, or course modules. However, less than half of interviewees reported having made use of university-provided resources, with interviewees from the humanities more likely to have done so than those in other disciplines.

It is important to note that when asked about using university provided resources, most interviewees interpreted this in terms of workshops, rather than online resources such as syllabi options, which were commonly used and appreciated. Interviewees who were less experienced with generative AI tended to find the workshops useful, but those who were more experienced tended to find them too basic and were more interested in discipline or even tool-specific training. As reflected in this study, faculty familiarity with generative AI varies widely, presenting a challenge for universities trying to cater programming to all levels.

---

<sup>33</sup> See Avi Staiman, “When Declarations Just Don’t Cut It: Building a Risk-Based Framework for AI Guidelines in Publishing,” *Science Editor* 48 (2025), <https://doi.org/10.36591/SE-4801-05>; Stephanie M. Lee, “Scholars Are Supposed to Say When They Use AI. Do They?” *The Chronicle of Higher Education*, December 18, 2024, <https://www.chronicle.com/article/scholars-are-supposed-to-say-when-they-use-ai-do-they>.

In some cases, interviewees were not aware of the resources their university had made available—whether workshops or even access to generative AI tools—until the interviewer informed them otherwise. An important task for higher education institutions is not only continuing to create support resources but also to promote available resources to their communities. It is also worth underlining that there was a significant uptick in support offered by institutions during the 2024-2025 academic year—through workshops, course modules, communities of practice, and more. It is likely that faculty awareness of and levels of participation in university-provided support has increased since our data was collected and will continue to do so as support opportunities on campuses proliferate.

The two most popular forms of support for instructors and researchers across disciplines were: 1) self-directed learning from online sources, and 2) learning from peers in their field, in formal or informal formats. When it came to self-directed learning, many interviewees turned to the web, as one associate professor in law described, “I’ve used the internet to teach me.” YouTube tutorials were a popular source of information about generative AI, as were internet forums and social media. Some interviewees were relying on online modules offered by other organizations or tech companies or had simply scoured the internet for resources in varied locations. Interviewees had high satisfaction with this type of self-directed learning because it allowed them to find resources tailored to their specific needs and do so on their own time.

Other than self-directed learning, interviewees most commonly reported learning from peers in the field. One professor in physics, for instance, when asked what resources they were using to navigate generative AI, responded with a very typical answer: “not very many formal things, conversations with colleagues about what’s been working in their research groups, a few panel discussions in which people have discussed what they are doing and what’s working and what some of the pitfalls are.” The professor went on to clarify that these conversations were happening everywhere from group lunches at their institution to casual chats at conferences. Interviewees highly valued guidance from trustworthy peers in their field who had already made headway in figuring out how to adopt AI. Resources coming out of professional or scholarly associations, therefore vetted by their peers, were also highly valued by interviewees.

Interviewees highly valued guidance from trustworthy peers in their field who had already made headway in figuring out how to adopt AI.

Interviewees who reported not having looked for support were usually non-adopters and their reasoning for not seeking support was the same as their reasons for non-adoption: they have not yet made the time, feel AI is not suited to core parts of their work, or have ethical concerns. Having not yet made the time to “deal with” the generative AI issue was the most common reason given, and such interviewees usually expressed plans to do so soon.

### **Support Resources Desired**

Instructors and researchers are eager for their universities to offer support related to generative AI. Even those with higher levels of familiarity with generative AI applauded their institutions for making an effort to educate their less experienced colleagues. When asked which unit they would like to see offering support resources, interviewees rarely had a preference. Their main interest was in seeing enough variance in timing and format (e.g., in-person and virtual, asynchronous and synchronous) to make events easy to engage with. When it came to what type of support they would like to see, interviewees most often expressed interest in resources that fell into the following categories: support tailored to specific disciplinary contexts, peer-to-peer learning opportunities, vetted information on generative AI and the product landscape, and secure access to generative AI tools for themselves and their students.

### **Discipline-Specific Resources**

Interviewees across disciplines thought that how and when generative AI could be appropriately applied in teaching and research would vary between disciplines, even if there were baseline commonalities. As a result, they also felt guidance and support would also have to be tailored to disciplines.<sup>34</sup> As one professor in engineering put it, “I think we should

---

<sup>34</sup> When asked what support or research would be helpful in incorporating generative AI into their research, approximately 90 percent of respondents to Ithaka S+R’s 2024 survey of biomedical researchers said discipline-specific support would be “slightly” to “very” helpful, in the form of either training or access to discipline-specific tools; see Dylan Ruediger, Chelsea McCracken, and Makala Skinner, “Adoption of Generative AI by

try to stay away from a one-size-fits-all approach. I think different programs are going to have very different needs. So, you can't just have one set of rules that are going to apply to everyone. So there needs to be enough flexibility.”

“So you can't just have one set of rules that are going to apply to everyone. So there needs to be enough flexibility.”

Nonetheless, most interviewees did still appreciate when their university provided base-level guidelines at a school or institution level but hoped these could leave enough flexibility for each discipline, instructor, or researcher to adapt the guidance as needed. To foster more discipline-centered initiatives, interviewees also wanted to see more organization and activity on the level of departments or related fields. Interviewees thought their institutions at an administrative level could do more to encourage or incentivize these smaller scale initiatives.

The strong interest in discipline-specific training often came from individuals with higher levels of familiarity with generative AI, who felt they were beyond the level of workshops being offered on AI basics. A key challenge for institutions is catering to an audience with widely varied levels of knowledge about AI to support beginning, intermediate, and advanced users. Institutions should also consider how to balance resources that are applicable institution-wide with other initiatives to build out more context-specific support and guidance.

It is also worth mentioning that, in contrast to claims from interviewees that they needed more discipline-specific support, our findings in this study suggest that the principles and ethics behind AI use are not always discipline specific. However, what may be more particular to a discipline are the specific templates or examples for how to apply generative AI. As will be elaborated on in the following section, one of the reasons interviewees are so eager to learn from their peers is to learn about concrete examples of AI use that they can subsequently try out in their own work. Determining where discipline matters when it comes to support

---

Academic Biomedical Researchers,” *Ithaka S+R*, October 17, 2024, <https://doi.org/10.18665/sr.321415>.



and applications for generative AI will be an important question for institutions and support units to keep in mind.

## **Peer-to-Peer Learning**

Opportunities to learn with or from one's peers were in high demand among interviewees. This included department-level conversations, presentations from colleagues with higher levels of familiarity with generative AI, and communities of practice. There was particular interest in longer-term formats, where a group would meet across multiple sessions, to foster more in-depth learning. As an associate professor in education explained, "I would love to be part of a semester-long or year-long professional learning community where we all have autonomy to use AI in different ways. Then we come together and we share it and learn from each other... that long term professional development, the continual professional development where you keep growing and your goals... I think they're so powerful." For a business professor, the group format was also appealing because it would hold them accountable to learning about generative AI. As they explained, "For me to learn something new, I need to be held accountable for it... because if I'm left to my own, obviously I'm not doing anything, right?"

Instructors and researchers trust their peers, hence their desire to learn from and with them. As one professor in health sciences explained, "When one faculty member says, oh, I'd use this in the classroom, or I've used this in research.. that kind of gives it a stamp of credibility where they've already done some of the legwork, and that makes me more likely to try it out." That "stamp of credibility" was particularly important for those, like one instructor in film, who defined themselves as "slow adopters." As they explained, "I'm not yet convinced that I need it, and so it's still just sort of learning more about it and kind of hearing how other instructors end up using it... That's how I'm going to learn whether it's something I really want to accept or reject."

It was important, too, for instructors and researchers to feel that they walked out of the learning experience with concrete examples of how to directly apply AI in their work. As a business professor noted, "at the end of it, I've got something that I can use in my class, right? Some product, some assignment, some improvement to my syllabus, some improvement to my exams, or whatever, that I can then implement." This made learning from colleagues within one's discipline appealing: see what a trusted peer has done in a similar course or research context, then replicate it in your



own work. Making the time to discover productive applications of AI was one of the major hurdles to adoption in teaching and research. However, if a trusted colleague had already—in the words of the aforementioned health science professor—done the “legwork” to figure out how AI can be effectively applied in teaching or research, this prevented instructors and researchers from feeling they needed to reinvent the wheel in figuring out best practices for their specific work.

## **Product Knowledge and Access**

Interviewees displayed limited knowledge of the generative AI product landscape for higher education. Across disciplines, interviewees most frequently mentioned ChatGPT.<sup>35</sup> Some interviewees bemoaned ChatGPT hallucinating sources or information, but did not mention the option of turning to tools grounded in vetted, reliable content. In one example, the interviewer suggested to their interviewee, a professor in political science, that they consider tools from Elicit, Consensus, or JSTOR. The interviewee was receptive to the suggestion but had not previously known about these tools. Other large language models, such as Copilot, Gemini, and Claude, were occasionally mentioned, as were a few products more specific to a discipline or use case, such as Adobe, Wolfram Alpha, or Github Copilot. Nonetheless, the dominance of ChatGPT and similar generalist models across the interviews indicates that instructors and researchers need to be better educated on the higher education-specific product landscape.<sup>36</sup>

Interviewees who had some knowledge of tools beyond ChatGPT reported not keeping up with new generative AI tools and features because they felt overwhelmed by the sheer number. These interviewees often expressed interest in the university creating a centralized, iterative resource with generative AI-related information and a list of recommended products. As an instructor in the arts put it, “curation” of all the information out there, including recommendations on tools, “would be really helpful because right now it’s the wild west. There’s a tool for everything, and as a teacher and professional, I struggled to figure it out.” As revealed by Ithaka S+R’s Product Tracker, keeping up with the generative AI landscape for higher

---

<sup>35</sup> Wiley’s ExplanAltions study of researchers similarly found that ChatGPT was, by far, the common tool for researchers to have heard of, and that awareness of usage of other tools was low. See “ExplanAltions: An AI Study by Wiley,” Wiley, February 4, 2025, <https://www.wiley.com/en-us/ai-study/for-researchers>.

<sup>36</sup> Ithaka S+R’s Generative AI Product Tracker (<https://sr.ithaka.org/our-work/generative-ai-product-tracker>) is a resource aiming to facilitate knowledge of the generative AI product landscape for higher education.

education is challenging. New tools and features continually emerge, making creating an iterative resource on AI news and products a challenging task for each institution to undertake individually. However, institutions could still benefit from continuing to curate as much information on generative AI as is reasonable, as well as directing faculty and students towards vetted external resources.

Interviewees also noted the importance of affordable access to high quality tools in a secure environment. Enterprise-level access to large language models was a common request for institutions that did not already offer it.<sup>37</sup> Many interviewees thought secure access through their institution would foster further experimentation, thus helping them and their colleagues become more familiar with generative AI. They also thought that either access to secure tools or robust guidance on which tools to use would be crucial in allowing instructors, faculty, and students to reap the potential benefits of AI adoption. When it came to their students, interviewees regularly reported feeling concerned that their students did not understand the privacy risks associated with AI and would prefer to see them using university-approved tools where privacy was assured. Expecting students to pay was another issue: a professor in health sciences, for instance, explained that their program was considering fostering further adoption of generative AI within their curriculum, but cost was the major barrier. While instructors clearly liked the idea of their students using university-provided tools, whether students are actually choosing to use those tools rather than their private accounts will be important to monitor moving forward.

Access to tools and to computational resources for cutting-edge research were of particular concern for certain STEM and social science researchers, but not exclusively: as a professor in classics pointed out, “access to computing power is going to be key for absolutely every field of human inquiry full stop.” This professor was one of a few who expressed hope that higher education could find innovative solutions that would not leave them at the mercy of big tech. As they explained, “of course, most of this computing power is not even in public hands or university hands. So, it's also important that universities do develop their own models, because

---

<sup>37</sup> Half of chief technology officers report their institutions do not grant students institutional access to generative AI tools, according to a recent Inside Higher Ed survey. See Colleen Flaherty, “The Digital Divide: Student Generative AI Access,” *Inside Higher Ed*, April 21, 2025, <https://www.insidehighered.com/news/tech-innovation/artificial-intelligence/2025/04/21/half-colleges-dont-grant-students-access>.

otherwise we are in the hands of private enterprises with different priorities and less transparency.”

Better product knowledge and access is important for instructors and researchers as they experiment and familiarize themselves with tools, as well as effectively and responsibly implement them in classroom and research settings. Universities have begun investing in generative AI access for their communities, but the future costs of this evolving technology and the financial implications for higher education remain unclear. While interviewees wanted to see their universities invest, they also recognized the complexity of the situation. In sum, supporting access to and development of AI models will continue to be important for higher education, but this is already and will continue to be complex to navigate.

## **Conclusion: Universities’ Current and Future Responses**

Our study demonstrates that familiarity and adoption levels among instructors and researchers are varied but rising. Experimentation with generative AI is widespread, from those who are responding to keep up with their students, to those who are genuinely excited about how AI might positively transform teaching, learning, and research. In the time since these interviews were conducted, the number of individuals within higher education who are highly familiar with AI has only increased. New technologies are emerging—particularly in the realm of agentic AI, a term very few of our interviewees in spring 2024 referenced. Now, the crucial task is managing the transition from the phase of exploration to responsible, well-informed usage of this technology. It is also important that the now higher number of heavy users have opportunities to share techniques and best practices they have discovered with their less experienced peers.

Universities will play an important role in guiding their communities through this transition to a period of both increased and well-reflected AI use. When asked to evaluate their university's response to generative AI in spring 2024, most interviewees were reasonably satisfied. Ultimately, interviewee satisfaction with their institution's response was tied to seeing indications that their university was reacting proactively. Interviewees appreciated messaging from senior leadership and the existence of task forces, committees, and workshops, which made them feel, as one language professor put it, "like I am in good hands." They spoke particularly highly of their centers for teaching and learning and of the syllabi language options they had provided. Most interviewees felt, as one business professor remarked, that their university was "doing the best it could" in a challenging situation.

A minority of interviewees raised important issues worth highlighting: universities, and higher education more widely speaking, might want to ensure they are making well-thought-out decisions about how to react to AI, rather than succumbing to the hype and feeling the need to react in the same way as other institutions. One business professor articulated this point of view well, explaining that their university is:

just trying to keep up with what everyone else is doing. Because it's in fashion to just accept it and embrace it, then everyone should accept it and embrace it, cause that's where we're at. But I don't think that that's necessarily the right call... We're jumping on this train because everyone else is. And if we don't, we'll get left behind. But I think that sometimes it's like the tortoise and the hare; there's some advantages to taking things slower than everyone around you. And that's what I kind of wish—our university would take a step back and make sure that we want to move forward in the same way. And sometimes it's okay to stand out differently and not do those things.

This professor clarifies that they are not "one of those people stuck in their ways"—they are not against AI adoption altogether. Instead, they argue that there are people like themselves who "understand this technology, which is why I don't want to just free fall and go for it. I want to be more conscientious about it. I want to really think about these things."

In June 2024, a task force sponsored by the Association of Research Libraries (ARL) and the Coalition for Networked Information (CNI) developed a set of imagined scenarios of AI-influenced futures for the research and knowledge ecosystem. Each future scenario was informed by how much society would adopt and adapt to AI, as well as whether society would be intentional about how AI is adopted and adapted.<sup>38</sup> Our study's findings indicate that adoption levels are on the rise. But—as the above comments indicate—we would be well-served to not neglect being intentional about adoption.

Interviewees in our study, from those who were enthusiastic about AI to those who were more skeptical, showed a strong commitment to promoting excellence in teaching, learning, and research. To enable this, the higher education community at large will want to ensure they are making conscious, reflective choices as consumers of AI. This is doubtlessly challenging, as AI tools have come onto the market faster than we can learn about them. That said, sharing insights about best practices and ethical standards for AI across the higher education community is crucial to prevent the risks inherent in non-adoption as well as the risks of widespread adoption without sufficient intentionality. As AI literacy and familiarity levels continue to rise, it will be important to determine what problems AI can actually productively address and ensure it does not unnecessarily create new ones.

## Recommendations

### Universities

- Articulate a strategic vision for generative AI in collaboration with campus communities and clearly communicate the vision to faculty, students, and staff.
- Foster cross-institutional conversation and programing about generative AI, in particular to coordinate initiatives to boost AI literacy and data security awareness among students, faculty, and staff.
- Incentivize support related to AI on the levels of individual schools or departments, in addition to initiatives on an institutional level.

---

<sup>38</sup> “ARL/CNI AI Scenarios: AI-Influenced Futures,” *Association of Research Libraries, Coalition for Networked Information, and Stratus Inc.*, June 2024, <https://doi.org/10.29242/report.aiscenarios2024>.

- Coordinate at an institutional level to build consensus on baseline standards for student use of AI that would apply across courses and majors.
- Provide secure AI environments for use by faculty and students and designate relevant staff to monitor changes in terms and conditions of licensed AI software.

## **Libraries**

- Leverage existing expertise in information and data literacy to establish robust programming for AI literacy instruction.
- Expand scholarly communication staffing and programming to help researchers ethically use generative AI and effectively communicate their use of generative AI.
- Hold training sessions to promote effective use of new AI search and discovery applications by students and faculty.
- Host conversations with disciplinary communities about the long-term implications of AI-mediated interactions with the scholarly record, and new ways of interacting with the scholarly record.
- Develop programming and resources to help faculty understand the IP and copyright issues associated with generative AI.
- Build consensus within library and archival communities about when and how to preserve and cite GAI generated outputs and inputs.

## **Centers for Teaching and Learning**

- Identify heavy users of AI among the faculty and facilitate opportunities for them to share applications and best practices with their peers.
- Emphasize to faculty the necessity of gaining familiarity and literacy about AI, while also encouraging critical reflection and intentional AI pedagogy.
- Develop programming to help instructors build assignments that facilitate the use of AI to support traditional disciplinary learning outcomes.
- Incentivize faculty to conduct scholarship on teaching and learning with AI to create institution-specific data for decision making about effective AI pedagogy.

## **University IT**

- To ensure technology decisions remain tightly tied to the university's core missions, engage in regular conversation with faculty, staff, and students on campus regarding their generative AI needs and practices.
- Ensure that students, faculty, and staff will have secure, affordable access to generative AI. If access is already provided, evaluate the degree to which it has been leveraged and barriers to further adoption.

## **Publishers and Funders**

- Prioritize developing more robust policies and nuanced vocabularies for generative AI use and disclosure in research and scholarly communication. Seek opportunities to build consensus with peer organizations and scholarly communities to promote baseline best practices and consistent terminology when possible.
- Support long-term research in the scholarship of teaching and learning to provide the empirical evidence necessary to make data-driven decisions about AI and pedagogy.



# Appendix A

The team members from these 19 institutions completed interviews at their respective institutions.

Institution	Team Members
Bryant University	Dave Gannon, Terri Hasseler, Suhong Li, Phil Lombardi, Allison Papini, ML Tlachac
Carnegie Mellon University	Lauren Herckis, Haoyong Lan
Concordia University	Mike Barcomb, Dianne Cmor, Ann-Louise Davidson, John Paul Foxe, Fenwick Mckelvey
Duke University	Linda Daniel, Yakut Gazi, John Little, Grey Reavis, Joe Salem, Xinzhu Wang
East Carolina University	Wendy Creasey, Jan Lewis, Ken Luterbach John Southworth
McMaster University	Erin Aspenlieder, Matheus Grasselli, Helen Kula, Kimberly Mason, Stephanie Verkoeyen
Princeton University	Sami Kahn, Zachary Painter, James Van Wyck, Anuradha Vedantham
Queen's University	Johanna Amos, Yasmine Djerbal, Lindsay Heggie, Selina Idlas, Angelique Roy, Nasser Saleh, Gavan Watson
Stony Brook University	Amanda Alicea, Peter Diplock, John MR Fitzgerald, Mona Ramonetti, Rose Tirotta-Esposito, Steven Wong
Temple University	Stephanie Fiore, Rachael Groner, Joe Lucia, Lori Salem, Nancy Turner
Wesleyan University	Rachael Barlow, Kevin Butler, Jeffrey Goetz, Amin Gonzalez, Mary Alice Haddad, Laura Patey, Rachel Schnepfer, Lauren Silber, Lynne Stahl, Khai Tran, Andrew White
Yale University	Lauren Di Monte, Alfred Guy, Julie McGurk, Kassie Tucker, Ryan Wepler
University of Arizona	Angela Cruze, Chris Griffin, Cas Laskowski, Maliaca Oxnam, Kristina Riemer
University of Baltimore	David Kelly, Jessica Stansbury, Nima Zahadat, Kevin Wynne

University of Chicago	Gillie Abdiraxman-Issa, Lynn Barnett, David Bietila, Scott Campbell, Taylor Faires, Robin Paige, Dina Ibrahim Rashed, Torsten Reimer, Elena Zinchenko
University of Connecticut	Xinnian Chen, Tom Deans, Sue Huang, Maryam Mageed, Laurie McCarty, Jailyn Murphy, Tom Scheinfeldt, Laurie Taylor
University of Delaware	Meg Grotti, Kevin Guidry, Erin Sicuranza, Josh Wilson
University of New Mexico	Robyn Gleasner, Laura Hall, Cree Myers, Todd Quinn, Jet Saengngoen
University of North Texas	Benjamin Brand, Yunhe Feng, Regina Kaplan-Rakowski, Sue Parks

## Appendix B

Tables 1 and 2 show the standardized disciplinary affiliations and ranks of all of the interviewees and of the sample.

**Table 1**

Discipline	Total #	Total %	Sample #	Sample %
Administration	4	2%	0	0%
Anthropology/Geography/Economics	6	2%	1	2%
Biology/Chemistry/Environmental Science	18	7%	3	7%
Business	22	9%	4	9%
Computer Science	12	5%	2	4.5%
Education	12	4%	2	4.5%

Engineering	25	10%	5	11%
Fine Arts	19	8%	4	9%
Health Sciences/Medicine	23	9%	4	9%
History	6	2%	1	2%
Humanities (other)	19	8%	4	9%
Interdisciplinary College	6	2%	1	2%
Law	9	4%	2	4.5%
Library	3	1%	0	0%
Linguistics	4	2%	1	2%
Literature and Languages	18	7%	4	9%
Math/Physics/Astronomy	9	4%	2	4.5%
Political Science	9	4%	2	4.5%
Psychology/Neuroscience	7	3%	1	2%
Science/Technology	2	1%	0	0%
Social Science (other)	13	5%	2	4.5%
Total	246	99%*	45	100%

*\*Due to rounding, total percentages do not always equal 100*

**Table 2**

Rank	Total #	Total %	Sample #	Sample %
Assistant Professor	33	13%	6	13%
Associate Professor	53	22%	8	18%
Professor	68	28%	13	29%
Researcher	6	2%	1	2%
Lecturer/Professor of Instruction	30	12%	6	13%
Other Non-Tenure Track	28	11%	6	16%
Postdoc	3	1%	0	0%
Graduate Student	12	5%	2	4.5%
Librarian	3	1%	0	0%
Other	10	4%	2	4.5%
Total	246	99%*	45	100%

*\*Due to rounding, total percentages do not always equal 100*

# Appendix C

Below is the interview guide that cohort members used to conduct interviews at their institutions.

## Pre-Interview Introduction

Generative AI refers to technologies that can create original content such as text, code, and images based on patterns identified in training datasets.<sup>39</sup> Popular consumer tools such as ChatGPT have made this technology widely accessible, and the use of Generative AI technology is rapidly transforming workplaces across sectors, including in higher education. As AI use becomes ubiquitous, universities need to understand how the technology is being adopted by faculty and students in order to assess how it can be harnessed effectively in support of teaching, learning, and research.

Within this context, [*name of institution*] is participating in a multi-institutional study to better understand instructional and research practices that make use of Generative AI. The following interview questions aim to help us get a better picture of how these technologies are impacting teaching, learning, and research, as well as what kinds of support and policies should be put in place moving forward. We will also share an anonymized transcript of this interview (and all other interviews conducted for this project) with Ithaka S+R, a not-for-profit research organization, who will use them to develop national findings and recommendations. We anticipate that the interview will take just under an hour.

Do you have any questions about the study and/or your participation before we get started?

Do you consent to this interview and to it being recorded?

---

<sup>39</sup> Adam Pasick, “Artificial Intelligence Glossary: Neural Networks and Other Terms Explained,” *New York Times*, March 27, 2023, <https://www.nytimes.com/article/ai-artificial-intelligence-glossary.html>.

# Interview Questions

## Introduction

1. How would you describe your level of familiarity and expertise with AI in general and with generative AI tools specifically?
2. In general, how have researchers in your field reacted to the advent of generative AI?

## Teaching and learning

1. Have generative AI tools made you think differently about how you approach teaching? How?
  - » *If yes, can you give me specific examples of how you've done so?*
    - Do you think your attempts were successful or not? Why?
  - » *If no, do you anticipate doing so in the future? Why or why not?*
2. Have you tried to incorporate generative AI tools into your instructional practices? *Examples: course development, assignment design, assessment, lectures.*
  - » *If yes, can you give me specific examples of how you've done so?*
    - Do you think your attempts were successful or not? Why?
  - » *If no, do you anticipate doing so in the future? Why or why not?*
3. How are you addressing the use of AI technology with your students? Are there tools or resources you have found to be most useful as you navigate your students' uses of AI technology?
4. What is the biggest challenge you've experienced when trying to integrate generative AI into your teaching?

Thanks for these responses. I'm going to switch gears now and ask a few questions about your research practices.

## Research

1. Have you experimented with incorporating generative AI or other AI tools into your research methods and workflow? *Examples: using generative AI to discover new primary or secondary sources, to synthesize scholarly literature, to brainstorm or outline, and to draft text.*
  - » *If yes, can you give me specific examples of how you've done so.*

- Do you consider those experiments successful or not? Why?
  - » *If no, do you anticipate doing so in the future? Why or why not?*
2. Have you experimented with using generative AI or other AI tools to prepare research outputs such as articles or presentations?
    - » *If yes, can you give me specific examples of how you've done so.*
    - Do you consider those experiments successful or not? Why?
    - » *If no, do you anticipate doing so in the future? Why or why not?*
  3. How is your field navigating the ethical implications of the technology? Are there any resources that you have found to be especially helpful within your discipline to navigate this issue?
  4. Are there any especially exciting or interesting uses of the technology that you've seen (or seen discussed) in your field?

Thanks for these responses. I'm going to switch gears now and ask a few questions about support needs.

#### Support needs

1. Have you made use of any training, tools, collaborations, or other resources in order to incorporate generative AI into your teaching and/or research?
  - » *Where did you find those resources? Examples: workshops offered by the Center for Teaching and Learning or Library, resources provided by scholarly societies, online tutorials.*
  - » *Where would you prefer these resources be made available to you moving forward?*
2. Looking toward the future and considering evolving trends in your field, what types of training or support will be most beneficial to researchers and/or teachers in your field?

I have just a few more general questions before we wrap things up.



## Conclusion

1. What has the university done (that you are aware of) in response to the rise of generative AI technologies?
  - » Are you satisfied with that response? What do you think the university could do to better support instructors and researchers moving forward?
2. Is there anything else you would like to share with us about generative AI in relation to teaching, research, and learning, that we have not already addressed?

Thank you for your time today. Our next step is to finish conducting interviews at [*institution name*] so that we can develop better capacities to support researchers and teachers at [*institution name*]. As I mentioned earlier, an anonymized transcript of this interview (and all other interviews conducted for this project) will be shared with Ithaka S+R, a not-for-profit research organization, who will use them to develop national findings and recommendations.

Do you have any final questions or concerns?