

DOI: https://doi.org/10.48009/3_iis_2026_125

Appropriate use? Exploring generative AI expectations across academic disciplines

Alana Platt, *University of Wisconsin-Whitewater*, platta@uww.edu

Avi Fleischer, *Oakton College*, mfleischer@oakton.edu

Abstract

This study explores how instructors across disciplines perceive and navigate the growing presence of generative AI in higher education classrooms. Through semi-structured interviews with instructors representing diverse academic disciplines, this study employed thematic analysis to examine how faculty members are integrating AI into teaching and learning, the challenges they encounter, and the evolving expectations surrounding appropriate student AI use. Although the study is limited by its sample size (N=9), findings suggest that instructors view AI as both a significant opportunity and a substantial pedagogical challenge. Participants identified potential benefits, including writing and programming support, personalized learning assistance, accessibility enhancements, and improved student engagement, while also expressing concerns regarding academic integrity, overreliance on AI, diminished student skill development, inaccurate outputs, and difficulties establishing clear classroom expectations. Participants emphasized that appropriate use is highly contextual and shaped by disciplinary norms, professional expectations, and ethical considerations.

Keywords: AI-Assisted Learning, Human–AI Collaboration, AI Literacy, Faculty Perceptions

Introduction

Higher education is rapidly adopting generative AI, but norms surrounding appropriate use remain fragmented and highly context-dependent. Instructors and students alike recognize the value of AI tools for both enhancing education and boosting productivity. Further, the industry expects the workforce to have AI competency. However, there is tension over the appropriate use of AI in higher education. These challenges are further complicated by the fact that higher education is not a monolith. AI-driven changes are occurring in widely differing courses and programs of study, with varying required skills and expectations. As a result, practices that may be considered appropriate AI use in one discipline or learning context may be viewed as problematic in another.

Some prior works have begun to look at the promise of AI in specific domains, such as writing support (Jin et al. 2025) and programming (Liang et al. 2024). What is appropriate behavior in one course or program of study may not be appropriate in others. For example, using AI for writing assistance may be allowable in a marketing class writing a report; however, it may be inappropriate to use it on an assignment for a business writing class. Expectations and acceptable uses of AI may also vary substantially across disciplines based on differing pedagogical goals, professional norms, accreditation requirements, ethical concerns, and

the potential consequences of errors. Programs in areas such as medicine, engineering, or other safety-critical fields may place greater emphasis on human oversight, verification, and maintaining a human in the loop, resulting in stricter expectations surrounding AI use. At the same time, AI tools may still provide value within each of these domains when applied appropriately and responsibly. Additionally, some types of tasks and problems may be more suitable for AI assistance than others, particularly those involving routine support, ideation, or low-risk activities, as opposed to tasks requiring substantial human judgment, domain expertise, or ethical decision-making. These differences also reflect varying conceptions of what constitutes learning, mastery, and appropriate professional preparation within different academic disciplines. As a result, practices that may be considered appropriate AI use in one discipline or learning context may be viewed as problematic in another.

While prior research has explored AI use within specific domains and learning contexts, less attention has been given to understanding how faculty across disciplines collectively perceive the opportunities, risks, and appropriate uses of generative AI in higher education. Existing work has largely focused on student outcomes or individual instructional settings, leaving a gap in understanding how disciplinary context shapes faculty expectations and perceptions of ethical AI use. Addressing this gap is important as institutions continue to develop policies, pedagogical strategies, and guidance for AI adoption across diverse academic programs.

Hence, we explore the following research questions:

R1: What commonalities exist among faculty across disciplines regarding the usefulness, opportunities, and concerns associated with generative AI in higher education?

R2: To what extent do faculty from different disciplines diverge in their perceptions of appropriate and ethical use of AI in higher education?

In this work, we aim to understand the current state of AI usage in higher education across disciplines. While other works have explored how it's used in specific contexts, this work expands on existing literature by identifying commonalities across disciplines. To this end, we conducted a series of semi-structured interviews with instructors across several domains. Utilizing thematic analysis, we extract and explore commonalities. In the next section, we present prior work on AI in higher education, as well as relevant work in learning psychology. The following section presents our methodology for collecting and analyzing the data, and then our results and limitations. We conclude with a discussion and areas for future inquiry.

Literature Review

In this section, we present relevant prior work on AI use in higher education. We start with a discussion of recent research leveraging AI to support student learning and outcomes. Next, we present works that explore the interaction between humans and AI in the learning process, followed by a discussion of relevant learning psychology literature. We then discuss how students use and perceive AI in workflows. Taken together, this provides an overview of relevant strands of research. This leads us to conclude that a gap in the literature exists around how AI perceptions converge or diverge across disciplines, which is discussed at the end of this section.

AI-Assisted Learning and Student Outcomes

Generative artificial intelligence (AI), particularly large language models (LLMs), is rapidly transforming teaching and learning practices across higher education. As these tools become more accessible, they are increasingly integrated into student workflows, raising important questions about their impact on learning. Further, assessment and instructional mechanisms are undergoing transformation as faculty respond to these rapid developments. Early research has identified both opportunities and challenges associated with AI-assisted learning, particularly in academic writing and knowledge work.

In recent years, literature has examined the impact of AI-assisted writing on student outcomes. For example, Cardon et al. (2023) demonstrate that AI-assisted instruction can improve writing quality, organization, and student motivation, particularly in English as a Foreign Language (EFL) contexts. However, they also identified concerns about contextual accuracy, overreliance on AI tools, and the long-term impact of AI-enabled education. Similarly, Kim et al. (2025) found that students perceive AI as serving multiple roles, including writing assistant, tutor, and peer collaborator. Overall, the results were promising, with findings such as a more productive writing process and improved writing performance. However, they noted both AI-related shortcomings (e.g., hallucinations, lack of human understanding) and student-related shortcomings (e.g., negative attitudes, lack of higher-order thinking). These studies suggest that AI has the potential to enhance both writing performance and student engagement, but there are barriers to effective use.

Human–AI Interaction and Learning Processes

Other work focuses on the processes through which students interact with AI tools. Nguyen et al. (2024) identify distinct interaction patterns, showing that students who engage in iterative, deep interactions with AI tend to produce higher-quality work than those who use AI in a linear or superficial manner. Similarly, Wang et al. (2024) find that students commonly use AI for brainstorming and idea generation, with varying degrees of reflection and self-monitoring. Together, these studies suggest that the effectiveness of AI-assisted writing depends less on access to the technology and more on how students engage with it during the learning process.

Related work further emphasizes the importance of understanding human–AI collaboration as dynamic and layered. Jin et al. (2025) propose that AI use can be categorized into levels, ranging from basic technical support to more advanced text transformation, with higher levels of engagement associated with improved outcomes such as critical thinking and motivation. Importantly, this work highlights the role of self-regulated learning strategies in shaping effective AI use.

Student perspectives further reinforce the complexity of AI-assisted learning. Studies indicate that students use AI in a variety of ways, including brainstorming, editing, and seeking feedback, and that their motivations evolve over time (Kim et al., 2025; Wang et al., 2024). While some students report improved efficiency and reduced anxiety, others express uncertainty about whether their skills are improving or whether they are becoming dependent on AI tools. These mixed perspectives highlight the importance of understanding AI use as a situated and context-dependent practice rather than a uniform behavior.

These differing interaction patterns suggest that the educational impact of AI may depend not only on the technology itself, but also on the learning processes students employ while using it. To better understand these dynamics, it is useful to consider insights from learning psychology. One relevant framework is self-regulated learning (SRL), which emphasizes learners' ability to plan, monitor, and evaluate their own learning processes (Zimmerman, 2002; Pintrich, 2004). In the context of AI-assisted writing, students who engage in iterative prompting, reflection, and evaluation may be demonstrating higher levels of self-regulation, while those who rely on AI for direct output may exhibit more surface-level strategies.

Learning Psychology, Scaffolding and AI-Supported Learning

Additionally, the concept of instructional scaffolding provides a useful lens for understanding how AI can support learning. Scaffolding involves providing structured support to help learners accomplish tasks that they could not complete independently (Vygotsky, 1987). This concept is closely related to the Zone of Proximal Development (ZPD), which describes tasks learners can complete with guidance but not yet independently. Within this framework, AI tools may function as adaptive supports that help students progressively develop skills and understanding. This concept has been widely applied in educational settings and instructional design (Taber, 2018).

In digital learning environments, scaffolding has been extended to intelligent tutoring systems and adaptive technologies (VanLehn, 2011). AI tools that guide students through questioning, feedback, or step-by-step explanations may function as scaffolds, supporting learning while gradually promoting independence. They can further be used to help create an iterative process in learning where the student, along with the help of AI, makes changes and demonstrates where revisions may be needed and how to make them. This iterative process may also reflect aspects of exemplification theory, in which learners develop understanding through repeated exposure to examples and modeled revisions (Zillmann, 1999). Similarly, prior work in learning psychology suggests that learners possess limited attentional and cognitive resources (Kahneman, 1973; Sweller, 1988; Bjork & Bjork, 2011). In this context, AI tools may help reduce routine cognitive burdens associated with drafting, revision, and information retrieval, potentially allowing learners to devote greater attention to higher-order thinking and problem-solving. Further, AI-enabled educational environments may support more multimodal pedagogical approaches by enabling interaction through multiple formats such as quizzes, podcasts, and interactive learning materials (Fleischer & Fleischer, 2024).

Student AI Workflows and Perspectives

Beyond individual learning processes, recent work has also highlighted the importance of critical AI literacy. Wang et al. (2025) propose a multidimensional model that includes awareness of AI limitations, strategic interaction with AI tools, and critical evaluation of outputs. This perspective emphasizes that effective AI use requires not only technical skills but also the ability to interpret and critique AI-generated content. These findings suggest that students must develop not only the ability to use AI, but also the judgment to use it appropriately. Research in applied domains such as software development provides additional insight into AI use in knowledge work contexts. Liang et al. (2024) find that developers primarily value AI tools for efficiency gains, such as reducing keystrokes and accelerating task completion, while also reporting challenges related to controlling outputs and ensuring alignment with task requirements. Although this work is not focused on educational settings, it suggests that similar patterns may emerge among students, particularly in disciplines such as information systems and programming. Programming tasks require higher-order thinking and problem-solving, and students in this context may also be susceptible to some of the shortcomings noted when using generative AI for writing, such as overreliance on AI tools.

Research Gap

Taken together, the existing literature demonstrates that AI-assisted writing offers both significant opportunities and meaningful challenges. Prior research has established that AI can enhance writing performance and support various stages of the writing process, while also introducing risks related to overreliance, reduced cognitive engagement, and variability in student outcomes. However, much of this work has focused on student use and outcomes, with less attention to how faculty interpret, structure, and respond to AI use within their courses.

Further, much of the existing literature has examined AI use within specific instructional or disciplinary contexts, leaving a limited understanding of how perceptions of AI may converge or diverge across disciplines. Accordingly, there remains a need for research that examines AI integration from an instructional and institutional perspective. In particular, there is limited understanding of how faculty across disciplines are adapting teaching practices, designing assignments, and guiding students in the use of AI tools in real-world educational settings. Addressing this gap is essential for developing effective pedagogical strategies and aligning AI use with broader learning objectives in information systems and related fields.

Methodology

This study employed a qualitative research design using semi-structured interviews to explore how faculty are approaching the use of generative AI in teaching, learning, and assessment contexts. This work aims to understand both individual practices and broader patterns related to instructional adaptation, student use, and institutional response across a variety of disciplines. Participants were recruited through professional networks and institutional contacts, and were contacted using an IRB-approved solicitation email. No compensation was provided to participants. In total, the dataset includes the perspectives of 9 interviews from across the following disciplines: Economics (1), Engineering (1), Information Technology (1), Marketing (1), Nursing (4), and Psychology (1).

The semi-structured interviews were designed to gather insights on 7 key topic areas: Policies and Context; Challenges and Goals, Current Approaches, Student Experience, Future Needs, Student Success and Inclusion, and Adoption Considerations. These topics were developed to capture a broad picture of how AI is being used, along with possibilities for new approaches and known challenges across a variety of disciplines and in a range of use cases. Specifically, our aim is to elicit perspectives on AI use policies, instructional strategies, student behavior, and future needs. Additional follow-up or probing questions were asked as part of the semi-structured interview process. The IRB-approved semi-structured interview questions are shown below in Table 1.

The interviews were conducted in a private setting via video conferencing software. Participants consented to the interviews being recorded and transcribed. In one instance, a technical issue resulted in the loss of the audio recording for a group interview. The researcher took detailed, contemporaneous notes during the session, capturing key points and attributing them to individual participants. While these notes do not constitute verbatim transcripts, they were included as supplementary data to support and triangulate themes identified in the individual interviews.

Table 1. Semi-Structured Interview Questions

Question #	Question Topic	Question Text
1	Policies and Context	What are your institution's (and your department's or your own) current policies or guidelines around student use of AI in coursework?
2	Challenges and Goals	What are the biggest challenges you see in helping students use AI responsibly in their writing and coursework?
3	Current Approaches	How are faculty and staff currently guiding or teaching students about appropriate ways to use AI tools in their assignments?
4	Student Experience	How do students seem to view or approach AI in their academic work—do they see it more as a shortcut, a learning tool, or something else?

5	Future Needs	If you could design an ideal approach or resource to help students learn how to use AI effectively and ethically, what would it look like?
6	Student Success and Inclusion	How could tools or resources that support responsible AI use also help promote student success, particularly for students with learning differences or disabilities?
7	Adoption Considerations	What would make a new tool or program valuable for your institution in terms of both maintaining integrity and supporting student learning with AI?

Data were analyzed using a reflexive thematic analysis approach. After an initial review of the interview notes and a subset of the transcripts, an initial codebook was developed. We then iteratively refined the code book as additional interviews were coded. Particular attention was paid to areas of convergence and divergence, especially across disciplines. Emphasis was also placed on tensions and contradictions in inter- and intra-participant perspectives. Limitations include reliance on self-reported faculty perspectives, variation across institutional contexts, and reliance on a single data source based on notes rather than recorded transcripts. This approach is appropriate for capturing a snapshot in time, examining how generative AI is being used in higher education, potential areas for growth, and practical and ethical limitations.

Results

This study aims to understand individual practices and broader patterns of AI use in higher education across a variety of disciplines. The analysis of our semi-structured faculty interviews reveals themes in the promise and shortcomings of AI in higher education settings, which relates to our first research question:

R1: What commonalities exist among faculty across disciplines regarding the usefulness, opportunities, and concerns associated with generative AI in higher education?

These themes are discussed and presented below.

AI as Both a Shortcut and a Learning Amplifier

A strong recurring theme that was present in all interviews was the dual nature of AI as an education tool. Specifically, whether AI was a learning amplifier or a hindrance hinges on its appropriate use. Participants described these dual roles of AI: both a shortcut for assignments and a tool to support learning. On one hand, instructors reported widespread student use of AI to generate outputs with minimal engagement, particularly under time pressure. This behavior was often characterized as “shortcut” use, including copying and lightly modifying AI-generated responses. Respondents linked this behavior to shallower learning and lower quality output.

However, instructors emphasized that AI can support deeper learning when used appropriately. Some examples included debugging code, refining ideas, or exploring alternative approaches. Several instructors noted that students who engaged in sustained interaction with AI tools produced higher-quality work and demonstrated stronger conceptual understanding. This distinction suggests that the impact of AI on learning is not inherent to the technology itself, but contingent on patterns of use.

This dichotomy was illustrated by several examples. In a technical setting, an instructor expressed this tension as “Some [students] just want the result, others use it to debug and learn...” It was observed that

“Everything became kind of the same... less creativity,” and that there was “Better code, but less innovation.” In the context of solving business tasks, another instructor observed that “[Students] knowing when to use is really hard...” and that given that AI is “Good for getting from A to B, bad for training”, overreliant use can lead to poorer learning outcomes. However, modifying assignments to incorporate AI explicitly and document outcomes leads to better results. The same instructor requires students to submit chat logs, and reflected that it was helpful to see their thought processes and “... see where students get stuck...”

Shift from Product-Oriented to Process-Oriented Assessment

Another recurring theme was a shift in assessment. Instructors described updating or creating new assignments designed to place greater focus on the process by which an assignment deliverable is produced, rather than assessing the final deliverable solely. This supports teaching responsible, effective AI use, as well as documenting that appropriate use. Assignment adaptations included requiring submission of prompts or chat logs, incorporating AI process reflections, as well as exercises to compare human and AI-generated outputs.

There was variance in how pervasive and developed these adaptations were, but these changes reflect a shift in how instructors evaluate how students think and interact with tools, rather than solely what they produce. Further, these techniques underscore that new AI-enabled workflows are being incorporated into instruction. One instructor noted “We engage students in learning how to apply and utilize generative AI. [We] encourage students to utilize it.”

Absence of Shared Mental Models for Appropriate AI Use

The lack of clear, shared expectations for AI use was a recurring theme in interviews. As one Instructor stated, many times students “do not know where to draw the line”. This problem is especially acute when expectations differ across assignments, instructors, and academic units. This inconsistency contributes to confusion and may inadvertently encourage inappropriate use. This was identified as a gap in instruction and institutional policy. Multiple respondents emphasized the need for explicit guidance regarding when AI use is appropriate, how AI may be integrated into student workflows, and clarity on how to evaluate AI-enhanced student deliverables. Some respondents noted that faculty would also benefit from clearer frameworks for appropriate and effective AI use.

Uneven Impact of AI on Student Learning

Another strong theme that emerged was the uneven impact of AI on student learning. A number of factors were identified as contributing to this unequal benefit. For example, some students had more prior knowledge of AI, which allowed them to utilize it more effectively. For example, students with greater experience with AI were better able to interact with AI appropriately, leading to better outputs. Those less skilled exhibited more naive interactions, accepting AI output uncritically. Differences in behavior patterns also had an impact on student learning. Students prone to procrastination were more likely to use AI as a shortcut, whereas students who planned ahead tended to use it more appropriately.

Further, differences between free and paid tools influenced the depth of student interaction. One instructor noted that before switching to a process-oriented assessment, students with paid versions of AI tools were performing better on assessments. One instructor stated, “...when students were relying on free versions... what I found is that students that didn't have paid accounts wouldn't ask as many follow-up questions or the depth of the conversations was generally not as good. Those students with paid versions, because they don't have the same token restrictions... there was a difference.”

Additionally, AI was viewed as beneficial for non-native English speakers and students with learning differences, particularly through translation, simplification, and alternative content formats. One instructor observed, “It can make it more of an even playing field... someone who is a non-native speaker having tools that help them understand what's being asked.”

At the same time, participants noted that students who might benefit most from AI support do not always use it effectively. Taken together, this implies that AI can have both beneficial and negative impacts on equity in learning. Equal access to advanced models and instruction on appropriate use can yield better learning outcomes and help support individuals with learning differences and non-native English speakers. However, if any students lack access to the same advanced models and instruction as their peers, it can put them at a disadvantage.

Institutional and Instructional Misalignment

Just as students and instructors are grappling with the implications of general AI, so too are higher education institutions. Instructors described a lack of coordinated institutional response to generative AI. Our respondents identified numerous institutional challenges, including: inconsistent or absent policies across courses and programs; uneven faculty adoption and engagement; limited or ineffective training initiatives; and slow institutional response relative to technological change. A respondent observed, “I think higher ed needs to be doing a better job of staking our claim... rather than tech companies pushing it forward, and we just keep reacting.”

One participant described their institution’s approach as “evolutionary,” rather than deliberate or coordinated. There was also variation in AI policies across disciplines (e.g., prohibition in some programs versus active encouragement in others); however, due to our sample's small size ($n=9$), it is unclear if these represent broader discipline-level variability. Further, respondents identified the need for more systematic faculty development.

Changing Role of the Instructor

Another important theme that emerged is the evolution of instruction in the age of AI. Multiple instructors were self-reflective about how AI changes their role: specifically, in what ways it replicates current teaching approaches, ways it can be fused with existing practice for a better outcome, and areas where a human mind and connection are more supportive to students. To this end, instructors described increased emphasis on: guiding students in effective AI use; designing assignments that promote critical engagement, and interpreting student-AI interactions as evidence of learning.

In some cases, AI tools were used to augment instructional capacity (e.g., AI teaching assistants, automated feedback, or diagnostic insights derived from student interactions). For example, one instructor noted that they could “...create a trainer GPT... wouldn’t give answers, but guide steps”. This shift reflects a broader transition toward instructors acting as designers of learning environments and mediators of technology use.

Emerging Tradeoffs and Design Tensions

Instructors also reflected on the future of AI in the classroom, anticipating increasingly sophisticated AI-powered tools. Their insights highlighted several unresolved tensions associated with AI integration. Participants noted that while AI has the potential to improve efficiency and the quality of student work, it may also reduce the cognitive effort associated with learning. Similarly, although AI can support creativity and idea generation, some instructors expressed concern that widespread use may lead to more standardized or homogeneous responses. Faculty also recognized a tension between accessibility and academic integrity, observing that tools capable of improving inclusion and supporting diverse learners may simultaneously

introduce new opportunities for misuse. Finally, participants described a disconnect between the rapid pace of AI innovation and institutions' ability to develop the policies, infrastructure, and faculty support needed to effectively integrate these technologies into teaching and learning.

Respondents also suggested that AI may expose weaknesses in assignment design, noting that poorly structured tasks tend to produce equally superficial AI-assisted responses. These tensions suggest that effective integration of AI will require not only new tools, but also reconsideration of pedagogical design and assessment strategies.

Differences Among Disciplines

Although many common themes emerged across the various areas of study, we also sought to understand differences in faculty perceptions of AI across disciplines. Specifically, our second research question was:

R2: To what extent do faculty from different disciplines diverge in their perceptions of appropriate and ethical use of AI in higher education?

It was observed that many of the use cases for generative AI remained consistent across disciplines, as demonstrated in the previous sections. Within our sample, however, there was a preliminary indication that faculty from disciplines with greater exposure to human harm expressed greater hesitancy toward the use of AI in certain contexts. Specifically, among the nursing faculty interviews, participants discussed substantial variation in AI policies across disciplines and programs, ranging from a complete prohibition on AI use to active encouragement and integration into coursework as appropriate to the task. The faculty member from psychology was more explicit in laying out ethical limitations of working with human data, providing the example that "...grad students can't input patient data" into AI. While ethical limitations exist across disciplines, the respondents from nursing and psychology were unique from the other participants in raising concerns about the ethics of using sensitive human data. However, given the small and uneven sample, it is not possible to determine whether this pattern reflects a broader disciplinary difference or is simply an artifact of the participants included in this exploratory study. Future research with larger, more balanced samples across disciplines is needed to determine whether this represents a meaningful signal or statistical noise.

Limitations

This study brings together the perspectives of higher education faculty who are integrating AI into teaching and learning, the challenges they encounter, and the evolving expectations surrounding appropriate student AI use. Although a number of interesting recurring themes emerged, the work has limitations. One limitation is the sample size of the semi-structured interviews ($n=9$). This work was conceived as a small exploratory qualitative study, and a number of stable core themes and repeating conceptual patterns emerged, leading the researchers to conclude that the dataset provided sufficient depth to address the research questions. Nevertheless, the small sample size should be viewed as a significant limitation, and the findings should be interpreted as preliminary rather than representative of higher education faculty more broadly. As such, the results are intended to generate insights and identify areas for future investigation rather than support broad conclusions regarding disciplinary perspectives on AI use. However, it remains possible that additional interviews, particularly from disciplines or institutional contexts not represented in the current sample, may have revealed additional perspectives or nuances not captured in this study (e.g., instructors in arts disciplines). Additionally, because only a single participant represented most disciplines and four of the nine participants were drawn from nursing, the cross-disciplinary comparisons are structurally imbalanced. Consequently, observed disciplinary differences—such as the greater ethical hesitancy expressed by participants from nursing and psychology—should be interpreted cautiously, as they may reflect the composition of the sample rather than underlying disciplinary norms. Future research should

seek larger and more balanced samples across academic disciplines to better examine potential similarities and differences in AI perceptions and practices.

Additional limitations include reliance on self-reported faculty perceptions rather than direct observation of classroom practices or student behavior. Future studies could strengthen these findings by triangulating faculty interviews with classroom observations, course materials, institutional AI policies, and student perspectives. Further, the self-selecting nature of our recruitment process may overrepresent faculty who are more enthusiastic about AI. This self-selection bias may have contributed to the generally positive perceptions of AI observed across participants and may limit the transferability of findings related to AI usefulness, adoption, and institutional adaptation. Participant experiences were also shaped by differing institutional policies, disciplinary expectations, and levels of AI adoption, which may limit transferability across contexts. Moreover, the apparent convergence in perspectives across disciplines should be interpreted cautiously, as it may be influenced by the exploratory nature of the sample, participants' shared institutional context, early-stage adoption of generative AI technologies, and the self-selection of faculty already interested in AI, rather than reflecting true consensus across higher education. Furthermore, because generative AI technologies and institutional responses continue to evolve rapidly, faculty perceptions and instructional practices may shift over time. Longitudinal research examining how faculty attitudes and classroom practices evolve as AI technologies mature and institutional guidance becomes more established represents an important avenue for future work.

Discussion & Future Work

This study explored faculty perceptions of generative AI in higher education to identify common themes and areas of divergence across disciplines. A number of consistent themes emerged across interviews despite differences in discipline and institutional context, suggesting the answer to R1 is that significant commonalities exist regarding usefulness, opportunities, and concerns. Across disciplines, participants described AI as useful for brainstorming, feedback, accessibility support, efficiency, and instructional scaffolding, while also expressing concerns related to over-reliance, reduced critical thinking, academic integrity, and the need to redesign assessment practices. Participants also consistently emphasized the importance of process-oriented learning, critical evaluation of AI outputs, and explicit guidance regarding appropriate AI use. These recurring themes suggest that many of the core pedagogical and institutional tensions associated with generative AI are not isolated to specific disciplines, but instead reflect broader challenges facing higher education.

However, findings related to R2 indicate less discipline-related variation than we initially anticipated. The most significant divergence was in the contexts of nursing and psychology, specifically around ethical limitations surrounding sensitive human data, such as entering patient or clinical information into AI systems. Outside of these areas, however, faculty across disciplines expressed surprisingly similar views regarding AI use, student behavior, and instructional adaptation. One possible explanation for this convergence is that higher education remains in the relatively early stages of generative AI adoption. As AI tools continue to develop and become more embedded in higher education, greater divergence in practice may occur.

Bahreini et al. (2023) proposed a four-stage adoption model for AI in higher education that may provide a useful lens for examining how faculty adoption evolves over time. While the present exploratory study was not designed to formally evaluate faculty against this framework, the interview findings suggest that many participants may be transitioning from the Availability stage toward Acceptability, recognizing AI's usefulness while continuing to grapple with its pedagogical and ethical implications. Future longitudinal

work could explicitly apply and validate this framework to examine how faculty progress through the proposed stages, including the transition to Stage 4 Adaptability.

Another avenue for future work would be applying the lens of Unified Theory of Acceptance and Use of Technology (UTAUT) (Venkatesh et al. 2003) to interpret faculty perceptions of generative AI in educational contexts. Faculty frequently described AI as highly useful for improving efficiency, supporting learning, and enhancing instructional capacity, aligning with the concept of performance expectancy. Similarly, many faculty noted the ease and accessibility of current AI tools, particularly in relation to rapid student adoption, reflecting effort expectancy. The interviews also highlighted the importance of facilitating conditions, including institutional policies, faculty development opportunities, technological infrastructure, and access to appropriate AI tools. There are also interesting questions on the impact that social influence may play in AI adoption within higher education. Several faculty described a sense of pressure to keep pace with rapidly evolving technologies, shifting student expectations, and broader institutional conversations surrounding AI. In some cases, faculty expressed concern that avoiding AI entirely could leave students underprepared for emerging workplace expectations. This social and professional pressure may contribute to AI adoption even in situations where pedagogical and ethical concerns remain unresolved. Future research could further explore how institutional culture, disciplinary norms, peer expectations, and perceived professional relevance shape faculty willingness to adopt or resist generative AI technologies.

Acknowledgements

The authors used ChatGPT to assist with editing and refining the manuscript's written content.

References

- Bahreini, A. F., Basu, C., Madan, M., Platt, A., & Sankaranarayanan, B. (2023). A four-stage model of ChatGPT adoption: application in tourism, healthcare, and education industry. *Issues in Information Systems*, 24(3).
- Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. *Psychology and the real world: Essays illustrating fundamental contributions to society*, 2(59-68), 56-64.
- Cardon, Peter, et al. (2023). "The challenges and opportunities of AI-assisted writing: Developing AI literacy for the AI age." *Business and Professional Communication Quarterly* 86.3.
- Fleischer, A. & Fleischer, B. (2024). Multiple modality pedagogical approach based on high engagement for diverse learners. *Psychology & Psychological Research International Journal*, 9(1), 1-4. <https://doi.org/10.23880/pprij-16000383>
- Jin, Fangzhou, Chin-Hsi Lin, and Chun Lai. (2025). "Modeling AI-assisted writing: How self-regulated learning influences writing outcomes." *Computers in Human Behavior* 165.
- Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice-Hall
- Kim, Jinhee, et al. (2025). "Exploring students' perspectives on Generative AI-assisted academic writing." *Education and Information Technologies* 30.1.

Liang, Jenny T., Chenyang Yang, and Brad A. Myers. (2024). "A large-scale survey on the usability of ai programming assistants: Successes and challenges." *Proceedings of the 46th IEEE/ACM international conference on software engineering*.

Nguyen, Andy, et al.(2024). "Human-AI collaboration patterns in AI-assisted academic writing." *Studies in Higher Education* 49.5.

Pintrich, P. R. (2004). A conceptual framework for assessing motivation and self-regulated learning in college students. *Educational psychology review*, 16(4), 385-407.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive science*, 12(2), 257-285.

Taber, K. S. (2018). Scaffolding learning: Principles for effective teaching and the design of classroom resources. *Effective teaching and learning: Perspectives, strategies and implementation*, 1-43.

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational psychologist*, 46(4), 197-221.

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance of Information Technology: Toward A Unified View. *MIS Quarterly* 1 September 2003; 27 (3): 425–478.

Vygotsky, L. S. (1987). *The collected works of LS Vygotsky: The fundamentals of defectology (Vol. 2)*. Springer Science & Business Media.

Wang, Chaoran, Zixi Li, and Curtis Bonk. (2024). "Understanding self-directed learning in AI-Assisted writing: A mixed methods study of postsecondary learners." *Computers and Education: Artificial Intelligence*.

Wang, Chaoran, and Zhaozhe Wang. (2025). "Investigating L2 writers' critical AI literacy in AI-assisted writing: An APSE model." *Journal of Second Language Writing*.

Zillmann, D. (1999). Exemplification theory: Judging the whole by some of its parts. *Media psychology*, 1(1), 69-94.

Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory into practice*, 41(2), 64-70.