

Encouraging Agentic AI Use in Higher Education: A Repeated-Measures Study of Student Attitudes and Intentions

Dr. Traci Carte
tacart3@ilstu.edu

Dr. Mangolika Bhattacharya
Mbhatt1@ilstu.edu

Luke Oess
ltoess@ilstu.edu

Illinois State University
Normal, IL 61761 USA

Abstract

Artificial intelligence (AI) tools are increasingly embedded in higher education possibly leading to enhanced learning outcomes when used to support, rather than replace, student cognition. However, inappropriate use may undermine critical thinking and promote over-reliance. To better understand how students engage with AI and to encourage more effective use, this study examines whether a structured intervention can influence student attitudes and intentions toward AI. Drawing on the Technology Acceptance Model (TAM) and Expectation-Confirmation Theory, we conducted a repeated-measures investigation across two contexts: a low-stakes laboratory experiment and a higher-stakes field study. The intervention emphasized guided prompt engineering strategies designed to position AI as a learning support tool rather than a shortcut to solutions. In low-stakes conditions with limited task knowledge, AI use reinforced positive attitudes, consistent with expectation confirmation. In contrast, in higher-stakes conditions where students had greater task familiarity, AI use led to negative disconfirmation and declining attitudes. Across both settings, behavioral intentions remained relatively stable, suggesting greater resistance to short-term change. These findings highlight the importance of modeling purposefully, agentic AI use in educational settings.

Keywords: Artificial intelligence in education, Prompt engineering, TAM, Expectation Confirmation

Encouraging Agentic AI Use in Higher Education: A Repeated-Measures Study of Student Attitudes and Intentions

Traci Carte, Mangolika Bhattacharya, Luke Oess

1. INTRODUCTION

Artificial intelligence (AI) has rapidly become part of university teaching and learning and most university faculty have experienced the impact. We have likely seen first-hand that when AI is used appropriately the student has a valuable tool enhancing his/her education, but when used inappropriately, students turn in work that is not their own and they forfeit the actual learning.

In information technology (IT) education, studies document performance gains for students when AI is appropriately used. However, concerns about dependence are growing, with AI sometimes providing solutions rather than scaffolding reasoning (Fan et al., 2025). Helping students understand the cost of misusing AI (i.e., poor performance, and lost educational opportunities) is critical. Moreover, the introduction of AI is happening against a backdrop of the changing relationship between students and the university. As a result, independent AI use is propelled by ease of use and immediate results amid increasingly transactional views of schooling. These conditions can incentivize cognitive offloading and risk atrophying foundational knowledge and critical thinking (Burns et al., 2026; McMurtrie, 2024). This tension is best characterized as agentic versus non-agentic use. Agentic student use of AI is composed of: deciding what to ask AI, evaluating appropriate use, strategic use of AI, and reflection on what was used (Dai & Lai, 2025). Alternatively, non-agentic use is passive, uncritical, and outcome-oriented interaction that displaces rather than supports learning.

Grounded in the technology acceptance model (TAM) (Davis, 1989) as well as Expectation-Confirmation Model (Bhattacharjee, 2001), this study examines whether a brief intervention can alter student attitudes and intentions toward AI to support effective, agentic AI use. The intervention focused on leveraging LLMs as guided-learning tools rather than shortcuts. Specifically, we explored the use of carefully designed and structured prompts that encourage students to actively participate in problem-solving, understand the underlying logic, and incrementally develop their solutions instead of copy-pasting complete solutions.

Using a repeated measures design, we hypothesized that both attitudes and behavioral intentions would change from time 1 to time 2 at least in part driven by our intervention and experience with AI. Data were collected from two studies: a lab experiment as well as a field study. Our results in both studies ultimately demonstrated changes in attitudes toward AI resulting from this limited intervention but little impact on behavioral intentions. Our lab experiment produced an increase in attitudes over time while our field experiment produced a reduction in attitudes. Implications for our confirmation versus disconfirmation results are discussed.

This study argues that (and provides some validation for) higher education shifting attention from the development of increasingly agentic AI systems to the cultivation of agentic AI use, demonstrating how guided prompt engineering can help students retain control over learning while benefiting from AI assistance.

2. LITERATURE REVIEW

In the rapidly evolving field of AI, Large Language Models (LLMs) have demonstrated significant capabilities across a wide range of applications, including natural language processing, content generation, and increasingly, code assistance (Touvron et al., 2023). Models such as ChatGPT and LLaMA are now integrated into development environments, including modern Integrated Development Environments (IDEs) like Visual Studio Code, enabling students to receive real-time support during programming tasks. While these tools offer powerful assistance, they also introduce concerns in educational settings, particularly the risk of students relying on full code generation with minimal engagement or understanding.

Many researchers have focused on how AI is impacting student learning and how faculty might better direct outcomes with some calling for more specific instruction in prompt engineering to enhance outcomes and confidence when interacting with generative AI (Woo et al., 2024; Lee & Palmer, 2025). Multiple meta-analyses report improvements in learning performance, motivation, self-efficacy, and related perceptions when AI is integrated within structured, faculty-directed designs (e.g., Wu & Yu, 2024).

Despite growing evidence that AI can enhance student learning, prior research focuses mainly on technological capabilities or general adoption, with limited attention to how structured use, such as guided prompting, affects student attitudes and intention. It also remains unclear whether short-term exposure under different task conditions produces consistent changes. As higher education grapples with how AI adoption can facilitate student use of AI in supporting learning, we argue emphasis needs to be placed on *how* they are using it. In other words, how do we help students become agentic users of AI.

This study draws on the Technology Acceptance Model (TAM) to explain college students' behavioral intentions to use AI. TAM posits that individuals' acceptance of a technology is primarily driven by perceived usefulness (PU) and perceived ease of use (PEOU), which shape attitudes toward the technology and, ultimately, behavioral intentions to use it (Davis, 1989). This model has proven very effective in understanding drivers of and intentions to use technology (e.g., see Scherer, et al., 2019 meta-analysis). Within TAM, behavioral intention to use a technology is the most proximal determinant of actual usage and has been validated as a surrogate for use (Venkatesh et al, 2003). Behavior intentions are primarily driven by users' attitudes which are in turn impacted by perceived usefulness, and perceived ease of use (Davis, 1989; Li, 2023).

In this study, we have opted to measure only the most proximate measures of use - attitudes toward technology and behavioral intentions. While PEOU and PU are cognitive beliefs, attitudes represent the overall positive or negative evaluation (Davis, 1989). A growing body of research indicates that direct exposure to AI technologies can significantly influence users' attitudes. Exposure increases familiarity, reduces uncertainty, and potentially leads to more favorable evaluations of AI tools. Empirical evidence using quasi-experimental methods demonstrates that students who interact with AI systems exhibit significant changes in their perceptions and acceptance of these technologies (Zhang et al., 2025). These findings align with broader technology adoption research, which emphasizes the importance of experience in shaping attitudes (Li, 2023). Accordingly, interventions that provide structured opportunities to engage with generative AI are expected to produce measurable changes in attitudes. Therefore, we expect that students' attitudes toward AI will differ from pre-intervention to post-intervention periods.

H1: Students' attitudes toward AI will significantly differ from pre-intervention to post-intervention.

From a learning and experience perspective, interaction with AI systems allows students to update their attitudes, potentially reinforcing components of technology acceptance. As students gain hands-on experience, perceived risks may diminish while perceived benefits become more salient, leading to stronger intentions to incorporate AI into their academic tasks (Li, 2023; Zhang et al., 2025). Further, extensions of TAM suggest that behavioral intentions are not static but can evolve over time in response to new information and direct interaction with technology (Venkatesh & Davis, 2000). Longitudinal and extended TAM studies show that experience and continued use shape perceptions of usefulness and, in turn, strengthen usage intentions (Venkatesh & Davis, 2000; Li, 2023). In educational contexts, structured interventions—such as guided AI use—serve as a mechanism for activating these changes by aligning attitudes with concrete task performance, thereby strengthening intentions (Li, 2023; Zhang et al., 2025).

H2: Students' behavioral intentions to use AI will change from pre-intervention to post-intervention.

3. RESEARCH DESIGN

Two studies, a laboratory experiment and a field study, were conducted using pre- and post-survey designs. In both, we conducted an AI and prompt writing education, an intervention, and pre- and post-survey data collection. Study 1 was a lab experiment; participants were drawn largely at random and

engaged in an experiment for which they had no immediate preparation nor investment in the outcome. Study 2, however, involved students from a senior-level database class and had just completed coverage of SQL with completed assignments and grades returned. The assigned SQL task they performed for this study was counted as homework (i.e., that had an investment in the outcome and generally understood whether the AI content was accurate). These two different treatments might be labeled low-stakes AI use versus high stakes. The differences are summarized in **Table 1**.

3.1 Measurement

Attitudes toward AI and behavioral intention were measured using validated scales adapted from prior literature. Items and their source are provided in Table 4. The slight difference in items between study 1 and study 2 is indicated with a slashed entry of either a programming or database focus. In addition, basic demographic data were collected: age, gender, major, class rank. Further, we employed the instructional manipulation check technique (Oppenheimer, et al, 2009) to capture whether respondents were paying careful attention to the questions and their responses. All observations with inaccurate responses to this check were removed.

	Study 1 (low stakes)	Study 2 (high stakes)
Sample selection	Drawn at random across majors, status, experience	Drawn from a specific course
Task description	General programming task with no impact on scholarly attainment	SQL assignment, graded for database class
Task knowledge	Varying	Specific to class content and previous assignments
Time spent	20 minutes	30 minutes
Outcome measure	None	Graded for accuracy

Table 1: elements of study 1 versus study 2

3.2 Participants

Participants were recruited for study 1 by collaborating with student clubs to promote the event, asking faculty to promote in their classes, and visiting classrooms of interested faculty. Recruitment was limited to IT majors to control for variance associated with major (Xia, et al., 2025). In all, 57 students participated. The demographic makeup of our samples is consistent with the population from which they were drawn (IT majors at a Midwestern university). Table 2 shows survey completions and response rates.

	Pre-survey #responses/response rate	Post-survey #responses/response rate	Matched responses #responses/response rate
Study 1	49/86%	42/73.6%	41/71.9%
Study 2	31/54.5%	23/40.3%	18/31.6%

Table 2: Total usable responses for pre- and post-survey items are provided below.

3.3 Research Process

We began with a pre-survey, followed by an intervention that included an overview of AI and prompt engineering as well as an assignment to apply the prompt engineering techniques just discussed, followed by a post survey. The AI overview included an overview of AI to ensure consistent understanding among participants, comparison of differences between popular GenAI models used today, security concerns, and prompt engineering. The process of developing the presentation content included the combination of findings from academic and professional literature to develop an educational intervention on the use of AI. Definitions of concepts were gathered from Microsoft and Coursera. MS Co-Pilot was the AI used due to the university having a license for it. Students were instructed to use this because this provided more control over the data and questions shared with the AI platform.

3.3.1 Intervention

Before the assignment, students were given specific prompts they could use. We develop these prompts by investigating and categorizing prompt engineering strategies tailored for educational use, aligning

them with specific learning objectives and programming tasks. Our approach emphasized fostering student engagement, improving conceptual understanding, and reducing over-reliance on automated solutions. By integrating these prompt strategies into assigned workflows, we aimed to create a more responsible and effective use of LLMs in IT education, supporting both skill development and academic integrity. Sample prompts and their intentions are provided in Table 3.

In study 1, a simple programming assignment with a complex programming assignment alternative was given. The objective of both assignments was to write a program that read a text-based data file, detect the delimiter, extracted and normalized names, and displayed them in a grid. The complex program required handling duplicate names, sorting the names alphabetically, and printing the total number of names read. The goal was to gauge the effectiveness of using high-quality, human-first prompts to write code. The participants were given 20 minutes on the assignment. They were not required to complete it.

Learning Stage	Sample Prompt(s)	Purpose
Problem Understanding	<i>"I need to read a text file of names where the delimiter is unknown. What are common strategies to detect the delimiter reliably?"</i>	Helps students understand the problem requirements and identify possible solution approaches before coding.
Pseudocode Development	<i>"Show me pseudocode for a function that detects delimiters."</i>	Encourages students to develop high-level logic and solution structure without focusing on syntax.
Parsing and Data Cleaning	<i>"How to split lines using the detected delimiter and skip empty tokens?"</i>	Guides students in handling real-world input data and building robust parsing logic.
Output Formatting	<i>"Provide pseudocode for a function that prints five names maximum per line."</i>	Helps students interpret and implement assignment-specific output requirements.
Testing and Validation	<i>"What tests should I write to ensure my program handles tabs, pipes, commas, blank lines, and trailing delimiters?"</i>	Encourages verification of logic through test cases and edge-case analysis before implementation.
Error Handling and Robustness	<i>"How should I handle file not found or permission errors in a program?"</i>	Introduces defensive programming practices and reliability considerations.

Table 3: Prompt development

Study two followed a similar structure (i.e., AI coverage, guided prompts, pre -survey and post-survey). However, the content and prompts were modified for an SQL-based assignment. The assignment was a set of 5 queries. Students provided a data model to the AI and used prompts that asked for strategies for query writing as well as debugging but not solutions. Surveys were slightly modified for study 2 (see Table 4). A full class period was devoted to study 2. Approximately 30 minutes of which was spent doing the task. It is likely some students spent additional time outside of class to sufficiently complete the task for a grade.

Fundamentally, study 1 examines initial exposure effects under low-stakes conditions, while Study 2 evaluates the same intervention under higher-stakes, performance-relevant conditions.

4. ANALYSIS

Given the modest sample sizes, analyses were limited to models with few predictors and guided by theory. Effect sizes were emphasized over sole reliance on null-hypothesis significance testing.

To investigate construct validity and reliability, we stacked time 1 and time 2 data from study 1. Although such pooling ignores the nested structure of the data and associated assumptions, it remains useful for evaluating overall construct validity when level-specific distinctions are not the focus (Bolger & Laurenceau, 2013; McNeish, 2018). Exploratory factor analysis and reliability results are reported in Table 4. We used a cutoff in the EFA of .4 as suggested by Stevens (2012). Two items were dropped to

provide acceptable validity and reliability. The measure of attitude maintained 7 out of 8 items that were slight modifications of an existing scale (Rogers et al, 2023). The items measuring behavior intentions were more than slightly modified versions of Davis (1989) to better reference AI intentions. While 3 of the 4 items provided appropriate validity and reliability outcomes, the items themselves are not quite reflective nature (i.e., multiple aspects of behavioral intention are covered) and should be considered more of a composite measure. Because the present study employs repeated-measures ANOVA on observed composite scores rather than latent-variable structural modeling, the primary concern is the substantive interpretation of the composite rather than bias in structural path estimates (Hair et al., 2022).

4.1 Study 1 Attitudes

Repeated-measures analysis of variance (ANOVA) was conducted to examine changes in attitudes toward AI from Time 1 (pre-use) to Time 2 (post-use) using our sample of 41 matched responses. The analysis revealed a significant main effect of time, $F(1, 40) = 4.29, p = .045$, partial $\eta^2 = .097$, indicating that attitudes differed significantly across the two time periods. Inspection of descriptive statistics indicated that participants reported lower attitudes prior to using AI ($M = 31.00, SD = 6.15$) compared to after ($M = 32.22, SD = 5.59$), reflecting a modest increase in attitudes following exposure to AI. The effect size was small-to-moderate (partial $\eta^2 = .097$), suggesting that approximately 9.7% of the variance in attitudes was attributable to the effect of time. A paired-samples t -test confirmed that the change in attitudes was statistically significant, $t(40) = -2.07, p = .045$, with a small effect size (Cohen's $d = -0.32$). indicating that attitudes were significantly higher after use. These results indicate that changes in attitudes over time were significant and positive. H1 supported.

Items	Loadings N=80
General attitudes toward AI (Rogers et al 2023), Cronbach's alpha .737	
AI technology could be a useful tool for writing computer programs/databases (Att1)	0.759
Use of AI technology to complete a prog./database assignment could enhance my grade (Att2)	0.751
AI technology could hinder my ability to learn how to program for myself (Att3) R	Dropped
The use of AI technology to do a prog./database assignment is a form of cheating (Att4) R	0.660
Using AI technology probably violates some university ethical standards (Att5) R	0.604
I am confident that I can complete prog./database assignments without AI technology (Att6) R	0.463
AI technology could be used in the future when I need to program/use databases in my job (Att7)	0.636
AI technology is going to revolutionize how software/database technology is developed (Att8)	0.615
Behavior intention Derived from Davis 1989) Cronbach's alpha .700	
I plan to use AI to complete class assignments (BI1)	0.683
I will include AI output in assignments I turn in (BI2)	dropped
I plan to use AI to support my learning (BI3)	0.907
I understand I should not have AI do my work (BI4) R	0.674

Table 4: Items and their loadings

4.2 Study 1 Behavioral Intentions

Repeated-measures ANOVA revealed that the effect of time on behavioral intentions was not statistically significant, $F(1, 40) = 2.84, p = .100$, partial $\eta^2 = .07$. Although mean scores were slightly higher for time 2 ($M = 9.56$) compared to time 1 ($M = 8.22$). This may be related to the small sample size and a very small effect size ($\eta^2 = .07$). H2 was rejected.

4.3 Study 2 Attitudes

In all, only 18 matched responses were collected in study 2 making a formal analysis difficult. Power analysis indicates results from a sample of 19 across two time periods can only be reliably interpreted

if the effect size is large ($f = 0.40$) (Faul, et al., 2009). Repeated-measures ANOVA results revealed a large effect size ($\eta^2 = .41$) and a significant effect of time, $F(1, 17) = 11.88, p = .003$. Attitudes were higher at Time 1 ($M = 37.67, SD = 4.55$) compared to Time 2 ($M = 31.94, SD = 6.24$) indicating a substantial decrease in attitudes over time. Results were further confirmed by a paired-samples t-test showing a significant decrease in attitude, $t(17) = 3.45, p = .003$, with scores declining from Time 1 ($M = 37.67, SD = 4.55$) to Time 2 ($M = 31.94, SD = 6.24$). The effect size was large ($d = 0.81$), indicating a substantial change over time. H1 was supported for study 2, but the change in attitudes was in the opposite direction. Our results from study 1 suggested confirmation of expectations for AI use while study 2 provides some support for an expectation disconfirmation.

4.4 Study 2 Behavioral Intentions

Repeated measures ANOVA was not statistically significant nor was the effect size large, $F(1, 17) = 0.29, p = .596$, partial $\eta^2 = .017$. This test failed to meet the effect size requirement. As such, the insignificant result for the effect of time on behavioral intentions could not be interpreted. For study 2, H2 was not supported.

Finally, because prior research has suggested that demographic variables may influence attitudes and behavioral intentions toward technology, interaction terms for gender, age, and ethnicity were included in all four models (attitudes and behavioral intentions for study one and study two). The demographic controls produced no significant findings suggesting that our results over time did not significantly differ across demographic groups.

5. DISCUSSION

Our study aimed to add insight into how AI is incorporated into classrooms and how students make decisions about AI use. Repeated measures ANOVAs were conducted to examine changes in attitudes and behavioral intentions over time. Results revealed a significant increase in attitudes in study 1 and a substantial decrease in attitudes in study 2. Differences between these studies include technology exposure was over a somewhat longer duration (20 versus 30 minutes), and students were more invested in their use of AI (programming assignment that did not have to be completed versus an SQL assignment that was evaluated as part of their course grade).

While not the initial theory driving our study, a potentially useful lens for interpreting the results is expectation-disconfirmation theory (Bhattacharjee, 2001), which posits that users form initial expectations about a system that are subsequently confirmed or disconfirmed through actual use. It is expected that university students will approach their studies with high expectations of AI's ability to support their educational efforts. In study 1, students, upon experiencing guided interaction with AI focused on enhancing learning and confirmed their positive beliefs. This aligns with longitudinal TAM studies which show that attitudes toward technology become more positive over time as users gain experience (Venkatesh and Davis, 2000), i.e., expectation confirmation. However, these students had no real investment in the assignment, suggesting this expectations confirmation of attitude toward AI may be driven by perceived ease of use rather than actual usefulness.

Study 2 results, while cautioning against over-reliance given the small sample size, exhibit expectation disconfirmation in students' attitudes toward AI. Upon experiencing guided interaction with AI focused on enhancing learning, participants' expectations were likely more closely evaluated due to motivation to complete the assignment well. As a result, perceived usefulness was more specifically evaluated, and expectations may not have been fully met, leading to negative disconfirmation and a corresponding decrease in attitudes. Importantly, the magnitude of the effect suggests that this change is not trivial. The large effect size indicates that time—and thus the experience of using technology accounted for a substantial proportion of the variance in attitudes. This is consistent with previous longitudinal studies of AI use with a deep understanding of the assignment (Sun et al., 2024) wherein students' attitudes declined after AI use. Again, these results must be viewed with caution – the small sample size can cause results to be skewed by idiosyncrasies in the data.

In contrast, no significant change was observed in behavior, suggesting that changes in attitudes did not translate into corresponding behavioral adjustments. This pattern highlights a disconnect between evaluative and behavioral responses, indicating that behavior may be relatively stable even in the face

of declining attitudes. From a TAM perspective, attitudes are typically expected to influence behavioral intentions and subsequent use (Davis, 1989). However, prior research has shown that the relationship between attitudes and behavior can weaken in contexts where use is mandatory or constrained (Venkatesh & Davis, 2000). In such cases, external factors may override individual evaluations, resulting in stable behavior despite changing attitudes. One plausible explanation is that while participants revised their evaluations over time leading to less favorable attitudes, these attitudinal changes may not be sufficient to alter behavioral intentions, particularly if participants continued to perceive strong social expectations or maintained confidence in their ability to perform the behavior. For example, in an environment where AI use is being equated with future employment. In this sense, attitudes may be more sensitive to short-term influences, whereas intentions reflect a more stable integration of multiple cognitive and contextual factors. This pattern is consistent with research suggesting that intention formation is influenced not only by cognitive evaluations but also by habitual, contextual, and normative factors (Li, 2023; Zhang et al., 2025). Alternatively, behavior intention is rather narrowly focused on *whether* the respondent will continue to use the technology – our results suggest maybe a different perspective is *how* will they continue to use it. We conducted a post-hoc analysis of our study 1 results for item BI4 only, “I understand I should not have AI do my work,” due to its clear focus on more agentic use of AI. The results showed a significant effect of time, $F(1, 40) = 7.07$, $p = .011$, partial $\eta^2 = .15$. The item score increased from Time 1 ($M = 2.00$, $SD = 1.22$) to Time 2 ($M = 2.85$, $SD = 2.14$), indicating a significant improvement over time, suggesting a small move toward agentic use. Study 2 provided no such result (perhaps due to sample size).

Taken together, elements of our study inform decision making around agentic use as well as mechanism to potentially elicit confirmatory versus dis-confirmatory experiences with AI use. As with all studies, this study has limitations. First the sample is very small, results should be confirmed with a larger sample. Second, more variables from both TAM and Expectation Confirmation could be collected to confirm results as well as to fully understand the drivers of these results. Specifically, PEOU and PU may add additional explanation for the changes in attitudes. Future repeated measures designs might benefit from using the ECT constructs. Further, these results may be specific to IT students. Their high expectations (perhaps unreasonably high attitudes) may reflect a willingness to see all IT innovation as good, as such future studies should consider student major field of study (or professionals’ area of expertise).

6. CONCLUSION

AI can enrich learning when embedded in sound pedagogy but may undermine foundational skills when used without guardrails (Wu & Yu, 2024; Li, 2023; Zhai et al., 2024; Burns et al., 2026). In computer programming education, Large Language Models (LLMs) and prompt engineering offer significant potential to support personalized and adaptive learning. By leveraging LLMs’ ability to interpret and respond to natural language, educators can create interactive environments that guide students through problem-solving processes while steering them away from over-reliance on fully generated solutions that can harm outcomes by enabling cognitive offloading.

In AI-mediated higher education, student learning outcomes are shaped not only by access to technological tools but by the extent to which students exercise agency in directing, evaluating, and reflecting on their interactions with those tools. As higher education navigates this AI environment, the shift from AI agency to student agency is crucial. This study aims to promote responsible and effective integration of LLMs in programming education, shifting their role from solution generators to interactive learning companions.

Acknowledgment: The authors would like to thank Daniel Freburg, Hannah Reichensperger, Jenna Pullega, and Jasmine Dhamers for their help in this project.

7. REFERENCES

- Bhattacharjee, A. (2001). Understanding information systems continuance: An expectation-confirmation model. *MIS Quarterly*, 25(3), 351–370. <https://doi.org/10.2307/3250921>
- Bolger, N., & Laurenceau, J.-P. (2013). *Intensive longitudinal methods: An introduction to diary and experience sampling research*. Guilford Press.

- Burns, M., Winthrop, R., Luther, N., Venetis, E., & Karim, R. (2026). *A new direction for students in an AI world: Prosper, prepare, protect*. Brookings Institution.
- Dai, Y., & Lai, S. (2025). *A theoretical framework of student agency in AI-assisted learning: A grounded theory approach*. Oxford Intersections: Social Media in Society and Culture. <https://doi.org/10.48550/arXiv.2512.07143> [arxiv.org]
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Fan, G., Liu, D., Zhang, R., & Pan, L. (2025). The impact of AI-assisted pair programming. *International Journal of STEM Education*, 12, Article 16. <https://doi.org/10.1186/s40594-025-00537-3>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149–1160. <https://doi.org/10.3758/BRM.41.4.1149>
- Lee, D., & Palmer, E. (2025). Prompt engineering in higher education. *International Journal of Educational Technology in Higher Education*, 22, Article 7. <https://doi.org/10.1186/s41239-025-00503-7>
- Li, H. (2023). ChatGPT-based flipped learning. *Australasian Journal of Educational Technology*, 39(5), 40–58. <https://doi.org/10.14742/ajet.8923>
- McMurtrie, B. (2024). Customers in the classroom. *Chronicle of Higher Education*. Muthén, B. O. (1994). Multilevel covariance structure analysis. *Sociological Methods & Research*, 22(3), 376–398. <https://doi.org/10.1177/0049124194022003006>
- McNeish, D. (2018). Thanks coefficient alpha, we'll take it from here. *Psychological Methods*, 23(3), 412–433. <https://doi.org/10.1037/met0000144>
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks. *Journal of Experimental Social Psychology*, 45(4), 867–872. <https://doi.org/10.1016/j.jesp.2009.03.009>
- Rogers M., Hillberg, H., and Groves. C., (2024). Attitudes Towards the Use (and Misuse) of ChatGPT: A Preliminary Study. In Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024), <https://doi.org/10.1145/3626252.3630784>
- Scherer, R., Siddiq, F., & Tondeur, J. (2019). The technology acceptance model (TAM): A meta-analytic structural equation modeling approach to explaining teachers' adoption of digital technology in education. *Computers & Education*, 128, 13–35. <https://doi.org/10.1016/j.compedu.2018.09.009>
- Stevens, J. P. (2012). *Applied Multivariate Statistics for the Social Sciences* (5th ed.). Routledge.
- Sun, D., Zhang, Y., Chen, X., & Wang, Q. (2024). Would ChatGPT-facilitated programming mode impact college students' programming behaviors, performances, and perceptions? An empirical study. *International Journal of Educational Technology in Higher Education*, 21, Article 13. <https://doi.org/10.1186/s41239-024-00446-5>
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., & Lample, G. (2023). Llama: Open and efficient foundation language models (arXiv:2302.13971). arXiv.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>

- Venkatesh, V., & Davis, F. D. (2000). A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, 46(2), 186–204. <https://doi.org/10.1287/mnsc.46.2.186.11926>
- Woo, D. J., Wang, D., Yung, T., & Guo, K. (2024). *Effects of a prompt engineering intervention on undergraduate students' AI self-efficacy, AI knowledge and prompt engineering ability: A mixed methods study*. arXiv
- Wu, R., & Yu, Z. (2024). AI chatbots and learning outcomes. *British Journal of Educational Technology*, 55(1), 10–33. <https://doi.org/10.1111/bjet.13334>
- Xia, Q., Zhang, P., Huang, W., & Chiu, T. K. F. (2025). The impact of generative AI on university students' learning outcomes via Bloom's taxonomy: A meta-analysis and pattern mining approach. *Asia Pacific Journal of Education*. <https://doi.org/10.1080/02188791.2025.2530503>
- Zhai, C., Wibowo, S., & Li, L. D. (2024). Over-reliance on AI. *Smart Learning Environments*, 11, Article 28. <https://doi.org/10.1186/s40561-024-00316-7>
- Zhang, X., Hu, X., Sun, Y., Li, L., Deng, S., & Chen, X. (2025). Integrating AI literacy with the TPB–TAM framework to explore Chinese university students' adoption of generative AI. *Behavioral Sciences*, 15(10), 1398 <https://doi.org/10.3390/bs15101398>