Student Perceptions of Agentic AI in Computer Science Education

A Thesis Presented for the

Master of Science

Degree

The University of Tennessee, Knoxville

Margaret Grace Kelley

May 2026

© by Margaret Grace Kelley, 2026
All Rights Reserved.

Dedication...

I would like to thank ...

My parents, Sharon and John, for all of that homeschooled science, helping me read the rules, and all of the praying about it.

My husband, Alex. This one I don't even know where to begin. Thank you for always being by my side. Even when we both knew it would be fine but I freaked out anyway, and you comforted me every time. I can't believe I not only got to do this, but we got to do it together. Boop.

My sister, Sonaz. My other half. The seagull has finally reached the shore.

My dearest friends, Yaren and Maddy. Coord, the game, Toothless, 191802, quiet pickles, and ladies, it is finally blackhole time.

Acknowledgments

I would like to give a special thank you to my advisor, Dr. Scott Ruoti, and the USER Lab. Working in the lab for the past two and a half years has shaped so much of my academic career. Honestly, it is the main reason I am here writing this thesis at all. Being part of the USER Lab gave me the opportunity to grow as a researcher, ask better questions, and realize that graduate school was something I actually wanted to pursue rather than just some far-off idea.

Dr. Ruoti's guidance, patience, and constant support played a huge role in that. He pushed me to think more critically, trust my own work, and take research seriously, while still making the process feel possible even when I was convinced everything was on fire. The opportunities and encouragement I received through the lab are the reason I decided to continue on to my Master's degree, and I am incredibly grateful for that.

I also want to thank John Sadik, Farzin Gholamrezae, and Allan Xie for all of their help with survey preparation and development. Building and refining this project would have been much harder without their feedback and support. A special thank you to Farzin for also taking the time to proofread what felt like a million (bad) versions of this thesis and helping me make sure it all actually made sense. Additionally, for also making the process feel possible even when I was convinced everything was on fire.

Abstract

Agentic artificial intelligence (AI) and large language models (LLMs) such as ChatGPT and similar AI systems have quickly become common tools in computer science education. Students are now using these systems for a range of tasks, such as explaining concepts, debugging code, generating examples, and assisting with coursework. While these tools may provide meaningful support to students– these models also raise concerns about the appropriate role of AI in educational settings. Although existing discussions often focus on what LLMs can do or cannot do, less is understood about how computer science students perceive these tools and their role in the learning process.

This thesis examines how computer science students perceive and use agentic AI tools in their coursework and learning. This study draws on anonymous survey responses collected from students in computer science programs. Computer science programs were selected from the top 300 Universities and Community colleges in the United States. This study focuses on computer science students' views regarding the usefulness of AI, the effect on understanding course material, and their role in the academic and professional spaces. The survey includes questions about student use of LLMs in coursework, their beliefs on whether these tools improve learning, and their broader concerns about ethics, environmental impact, privacy, academic integrity, and educator involvement.

Table of Contents

1 Introduction **1**

1.1 The Rise of LLMs in Computing Education 1

1.2 Goals of this Thesis 2

1.3 Contributions 3

2 Related Works **4**

2.1 LLMs as Learning Aids in Computer Science 4

2.2 Ethical Considerations and Academic Integrity 5

2.3 Summary and Research Gap 6

3 Methodology **7**

3.1 Survey Design and Development 7

3.2 Participant Recruitment and Data Collection 8

3.3 Data Processing and Analysis 10

3.4 Participant Demographics 10

4 Results **14**

4.1 General AI Familiarity and Access 14

4.2 Usage Frequency 15

4.3 Perceived Usefulness 18

4.4 Trust and Reliability 21

4.5 Impact on Learning and Motivation 23

4.6 Perceptions of AI in Future Careers 24
4.7 Qualitative Findings 25
4.8 Summary of Results 27

5 Discussion **30**
5.1 What the Data Tells Us 30
5.2 The Trust Gap 31
5.3 Implications for Learning 32
5.4 The Emotional Dimension 33
5.5 Limitations 33
5.6 Recommendations 34
5.7 Future Research Directions 35

6 Conclusions **37**

Bibliography **39**

Vita **43**

List of Figures

3.1 Overview of the survey structure used to collect responses regarding student experiences with agentic AI tools. 9
3.2 Distribution of participants by university type. 12
3.3 Distribution of participant gender identities. 12
3.4 Distribution of completed years of computer science coursework among participants. 12
3.5 Distribution of self-reported GPA among participants. 13

4.1 Length of participant experience using agentic AI tools. 16
4.2 Reported paid use of agentic AI tools among participants. 16
4.3 Student perceptions of the extent to which agentic AI can fill various academic roles. 17
4.4 Reported frequency of agentic AI use across coursework-related tasks. 17
4.5 Reported frequency of agentic AI use across coursework-related tasks. 19
4.6 Student perceptions of the general usefulness of agentic AI tools. . . . 19
4.7 Student perceptions of the usefulness of agentic AI tools across coursework-related tasks. 22
4.8 Student confidence in the accuracy and reliability of agentic AI outputs. 22
4.9 Student perceptions of the role of agentic AI tools in future computer science careers. 29

List of Attachments

Recruitment Email (Request_for_Survey_Participation_Email.pdf)

Consent Form (Consent_for_Research_Participation.pdf)

Qualtrics Survey (Qualtrics_Survey_AAI.pdf)

Survey Data (Survey_Data.xlsx)

Chapter 1

Introduction

Large language models (LLMs) have moved quickly from research curiosity to an everyday tool [20, 17]. Students are now routinely turning to these systems for help writing code, debugging errors, and understanding course material [24]. In many cases, LLMs serve as the core component of broader agentic AI systems. For computer science students in particular, these tools are increasingly incorporated into their academic workflows [1, 21].

However, this rapid adoption has outpaced our understanding of what it actually means for students. That is to say, whether LLMs ultimately help or hurt student learning remains an open issue [17, 21]. Although prior research has explored how students use AI [21], [17], there is still a gap in understanding how students themselves *perceive* these tools, and how those perceptions shape the way they influence the learning process.

1.1 The Rise of LLMs in Computing Education

The growing presence of agentic AI in computer science classrooms has introduced both opportunities and concerns within computer science education [9]. On one hand, these systems provide on-demand assistance that many students have never

had before—particularly those who lack easy access to office hours, tutoring, or study groups. [11].

On the other hand, educators have raised legitimate concerns about what happens when AI-generated support becomes a crutch [15]. *Specifically, there is concern that over-reliance on these tools may undermine long-term retention and limit the development of fundamental skills* [16, 8]. If students offload that cognitive work to an AI, they may be trading short-term convenience for long-term deficits.

As AI continues to improve, these concerns become more significant. AI that can write a functioning homework submission in seconds poses a fundamentally different challenge than a search engine that just returns documentation. Understanding how students are using and thinking about these tools—not just whether they use them, but how they perceive them— is essential for designing policies that get the benefits without paying the costs [21, 14].

1.2 Goals of this Thesis

This thesis *examines how computer science students perceive and use agentic AI tools in their academic coursework.* The central themes are as follows: Do students think these tools are *actually* helping them learn, or just helping get things done? Does using AI make students more confident, or does it quietly erase their trust in their abilities? And what role, if any, do students think AI should play in their education going forward?

To investigate these questions, this study analyzes survey responses from students enrolled in computer science courses across the top 300 universities and community colleges in the United States. The survey examines students' familiarity with agentic AI tools, in what ways and how frequently they are used for coursework, and how students compare AI systems to other sources of assistance, such as professors, TAs, and tutors.

Participation in this study was voluntary and anonymous. All responses were collected through a Qualtrics survey, and participants were informed about the purpose of this research prior to participation.

1.3 Contributions

The findings from this work contribute to an ongoing discussion that is only going to grow more urgent as AI tools become more capable and embedded in academic life. By examining how students perceive AI—not just how they use it—*this thesis seeks to provide insight into how these tools are shaping the learning environment for computer science students.* [9], [11], [14].

This perspective matters for instructors when designing courses, departments when setting AI policies, and researchers building the next generation of educational technology. As these tools continue to evolve, understanding student experience is not a peripheral concern—*it is a prerequisite for responsible integration.*

Chapter 2

Related Works

The rapid integration of LLMs into computer science education has generated a growing body of research [20, 17, 21, 13], though much of it has concentrated on system capabilities and short-term performance outcomes. Less attention has been paid to the student experience itself—how learners perceive these tools, whether they trust them, and what role they believe AI should play in their education.

This chapter presents the study within the broader literature at large, displaying what is known about AI adoption in academic settings, the specific challenges that arise in programming contexts, and the ethical questions that have emerged alongside widespread use.

2.1 LLMs as Learning Aids in Computer Science

A steadily growing body of research has explored how LLM-powered tools are being integrated into computer science education. Prior work suggests that these systems have multiple benefits: improve student productivity, lower the barrier to on-demand academic support, and offer explanations and examples that supplement traditional resources [11, 20, 13, 25].

Students tend to use LLMs in ways that are additive rather than substitutive. They use LLMs to clarify lecture material, generate additional examples, or work through assignments when other resources are unavailable [21, 25].

That immediacy, though, cuts both ways. AI-generated outputs are not always correct, and incorrect responses are frequently delivered with the same confident tone as accurate outputs. Students—particularly those earlier in their education paths—may not have the knowledge to recognize when an answer is wrong. This introduces a quiet reliability problem that usage statistics cannot capture alone. More broadly, some researchers have raised concerns that frequent reliance on AI assistance may lead to over-reliance on AI-generated assistance. This in turn contributes to weaker development of the problem-solving skills that computer science education is specifically designed to build [16, 8, 22, 4].

2.2 Ethical Considerations and Academic Integrity

The increasing presence of agentic AI in academic settings has raised some ethical questions. At the center of these debates is the question of academic integrity; given that LLMs are capable of generating complete and plausible solutions to assignments, it is not always clear where the line falls between *acceptable assistance* and *academic dishonesty* [17, 21, 6].

Researchers and educators are not yet in agreement on the best use of AI tooling. Some argue that AI tools are analogous to other learning aids—such as calculators or Stack Overflow— and can be responsibly incorporated into coursework with appropriate guidelines [17]. Others express concern that unrestricted use creates opportunities for students to bypass the learning process [8]. This may lead to long-term consequences of which may not be visible until much later.

Beyond academic integrity there are broader concerns about trust and reliability. AI can generate responses that appear authoritative but contain factual errors or

hallucinated information [12, 19, 2]. For students without sufficient background knowledge to evaluate AI outputs critically, this is a problem [16, 8, 22, 20, 21].

Alongside these concerns, there are wider societal questions about bias, transparency, and environmental costs of large-scale AI deployment [7, 3]. While somewhat out of scope of this study, they form part of the broader context in which students are encountering these tools.

2.3 Summary and Research Gap

Taken together, prior work paints a picture of a technology that is genuinely useful, genuinely risky, and frankly—genuinely hard to study. Agentic AI tools have demonstrated real benefits in terms of accessibility, efficiency, and on-demand support. What literature still lacks is a closer look at the student perspective—not usage patterns or performance outputs, but *how students themselves understand and evaluate the role these tools play in their learning process.*

That gap is the focus of this study. By collecting survey data from students across a wide range of institutions, this research aims to provide a more grounded look at how computer science students perceive agentic AI; whether students find AI useful, if there is trust or lack thereof, how they think it affects their learning, and what role they believe it should play going forward. Understanding those perceptions is a necessary step toward designing educational environments where AI serves the student, rather than the other way around.

Chapter 3

Methodology

This study takes a mixed-method approach to examining how computer science students perceive and use agentic AI. Quantitative data—collected through structured Likert-scale and multiple choice questions— provides both measurable trends and more nuanced insight into student experiences. Qualitative data—collected through open-ended responses, expands on the trends. The reasoning, hesitation, and context that numbers alone can't capture. Together, they allow for a more complete picture of student experiences than either method would yield on its own.

Before survey distribution, all recruitment materials—the survey, invitation email, and consent statement, which can be found in the supplemental material document—were submitted for Institutional Review Board (IRB) approval.

3.1 Survey Design and Development

The development of this study began with an initial review of existing literature on AI in education to identify gaps and guide the formation of research questions. After reviewing prior work, a broad set of potential research questions was generated. This gap-mapping process shaped the research questions that would eventually anchor the survey instrument.

From there, a broad pool of roughly 50 candidate questions was generated and brought to my advisor, Dr. Scott Ruoti, for review. Questions were evaluated across multiple rounds of discussion—assessed for clarity, relevance, and alignment with the goals of the study. Peer researchers contributed additional feedback, and early versions of this work were presented at a EURēCA poster session, where additional feedback was collected and incorporated.

Once a working set of questions began taking shape, the survey was built in Qualtrics. Pilot testing with a small group of participants led to more feedback that wasn't visible in the design phase. These findings drove another round of revisions.

The final survey instrument consists of five primary sections, with a sixth section conditionally presented: participants who indicated no prior use of AI tools were not shown questions specific to AI in computer science contexts. The five core sections cover consent, general agentic AI usage, agentic AI usage in computer science, perceptions of agentic AI in CS contexts, and participant demographics. The complete survey is included in the supplemental material document. A complete view of the survey structure can be seen in Figure 3.1.

3.2 Participant Recruitment and Data Collection

Participants were recruited from 300 institutions across the United States. These institutions included a mix of four-year universities and community colleges, selected using the [5] database.

Once IRB approval was obtained, the survey was distributed through the institution-specific mailing lists, with each link encoded to track the school for later demographic analysis.

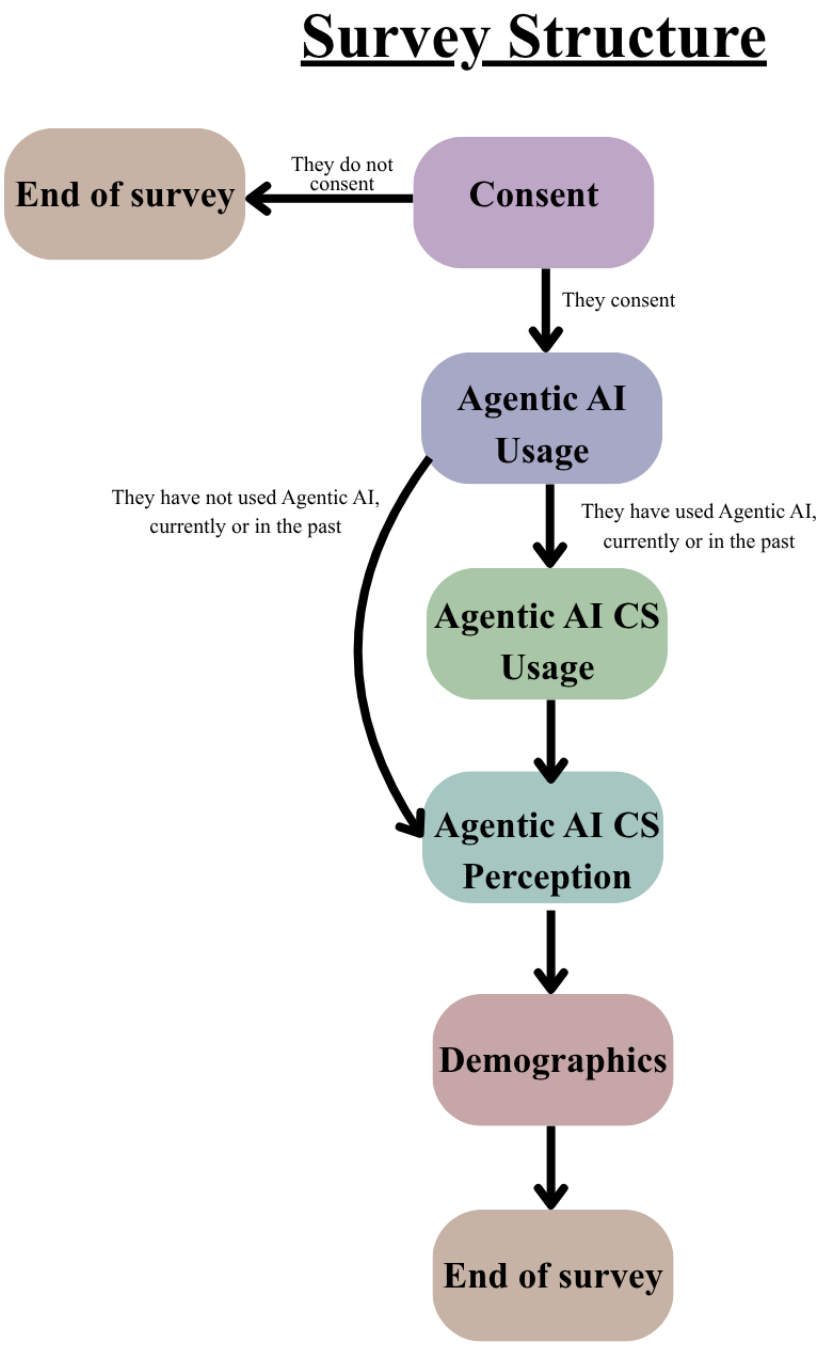


Figure 3.1: Overview of the survey structure used to collect responses regarding student experiences with agentic AI tools.

3.3 Data Processing and Analysis

Following the data collection, responses were cleaned and prepared for analysis. Incomplete, non-serious, or otherwise invalid responses were removed to ensure data quality.

The quantitative responses were standardized and prepared for visualization and statistical analysis. The qualitative responses required a different approach—each open-ended answer was manually reviewed and coded using Boolean indicators to mark the presence or absence of specific themes. This made it possible to integrate qualitative data into the broader analysis, rather than treating them as a separate, incomparable dataset.

Results were visualized using stacked horizontal bar charts throughout. This allows for normalized, consistent comparison across questions and response categories—making patterns visible in a way that raw counts would not show.

3.4 Participant Demographics

A total of 291 participants completed the survey, recruited from 290 institutions across the United States. Both universities and community colleges were selected from the top 300 schools in America listed in [5]. This sampling approach was intended to capture a diverse set of perspectives from students' experiences across different academic environments.

Demographic data was collected to better understand the composition of the participant group. This included institution type, gender, race, age, year in school, GPA, and both self-diagnosed and professionally diagnosed mental and physical disabilities.

The sample skews heavily toward university students: out of the total 291 participants, 3% (n=7) were identified as community college students, while 97% (n=284) were identified as university students, as shown in Figure 3.2. Despite this

imbalance, responses from all groups were included to provide a comprehensive view of student experiences. Excluding the community college cohort would reduce rather than improve representativeness—but findings should be read with that distribution in mind.

Participants were also asked to report their gender identity. Gender identity followed a pattern consistent with broader trends in the discipline. The majority of respondents identified as male 70.1% (n=204). This was followed by 22.3% (n=65) identifying as female. A smaller portion, 7.6% (n=22), selected "Prefer not to answer" or another category. Figure 3.3 shows the distribution of responses across gender identities.

This distribution is generally consistent with trends observed in computer science populations. [18] reports that in 2021, approximately 22.9% of computer science bachelor's degree graduates were women. A figure that closely mirrors this sample's composition, suggesting the participant pool reflects the field's existing demographic profile rather than a sampling artifact.

Participants represented a range of academic levels, from early undergraduates to graduate students—measured by years of completed computer science coursework, shown in Figure 3.4.

Academically, the sample trends toward stronger performance. The largest concentrations fell in the 3.7-4.0 and 3.0-3.7 GPA ranges, seen in Figure 3.5. This is worth noting when interpreting the results—students with higher GPAs may engage with AI tools differently than those with lower GPAs.

Finally, participants were asked whether they experienced any physical or mental health challenges affecting their academic work. The majority of participants reported no health-related challenges (45.3%, n=131). Mental health conditions were reported by 35.6% (n=103) of participants; physical conditions by 3.8% (n=11); and both by 15.2% (n=44). This distribution is included here as context for interpreting variation in student experiences—particularly in sections where qualitative responses touch on stress, workload, and academic pressure.

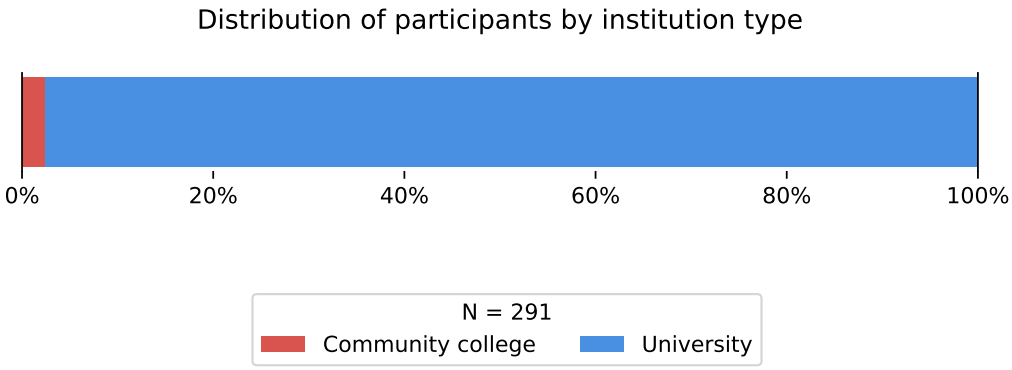


Figure 3.2: Distribution of participants by university type.

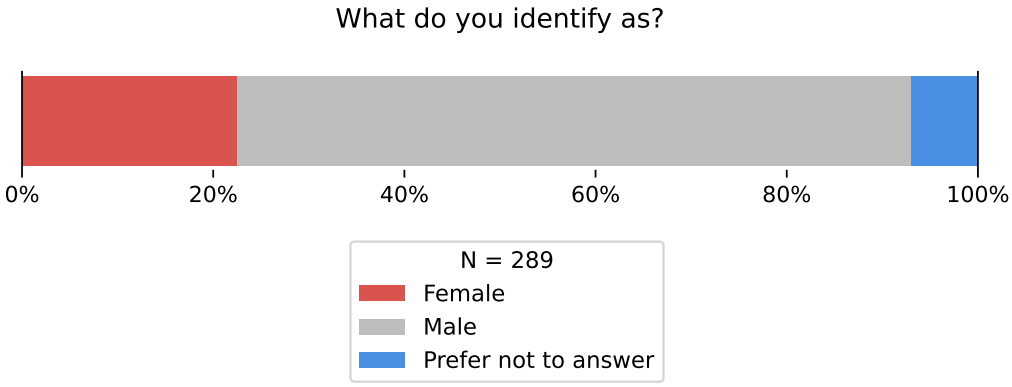


Figure 3.3: Distribution of participant gender identities.

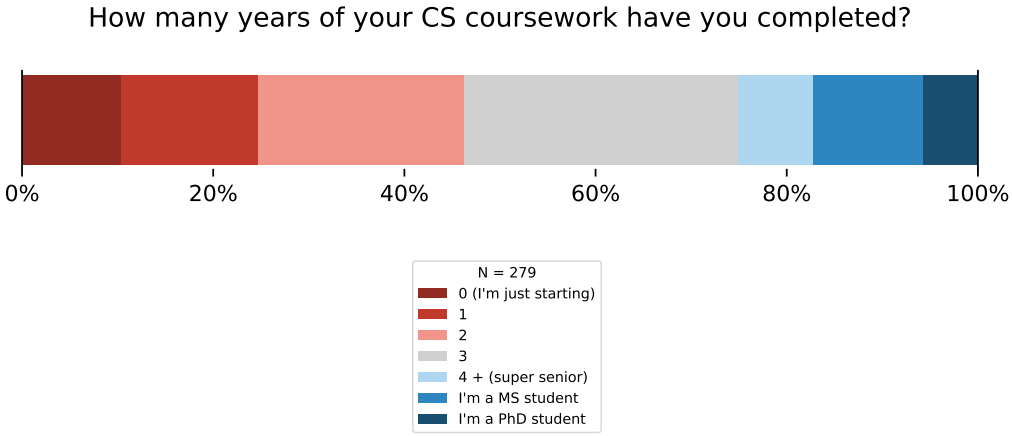


Figure 3.4: Distribution of completed years of computer science coursework among participants.

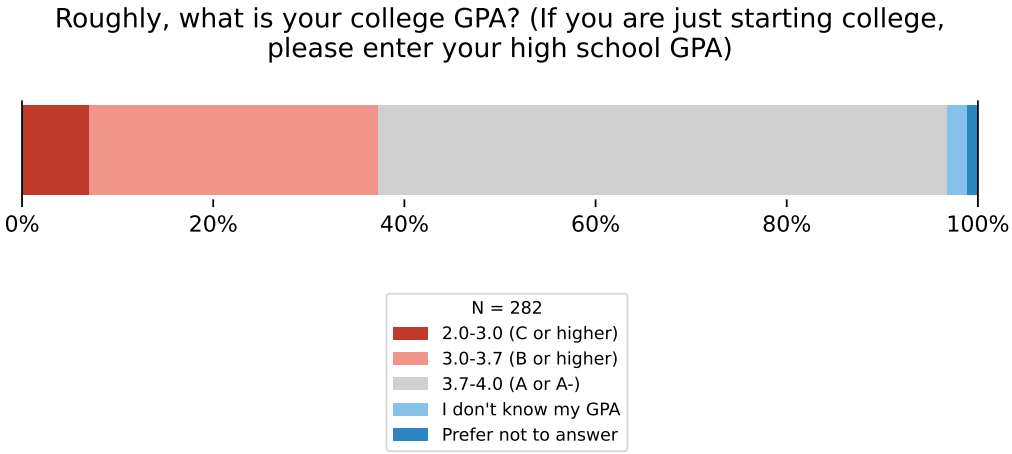


Figure 3.5: Distribution of self-reported GPA among participants.

Chapter 4

Results

This chapter presents the survey results on student perceptions and use of agentic AI tools. Results follow the structure of the Qualtrics survey, which can be found in the supplemental materials folder. Participants responded to Likert-scale and categorical questions; findings are reported using response distributions and percentages. Where appropriate, stacked bar charts are used to show the distribution of responses across categories.

4.1 General AI Familiarity and Access

Participants were asked about their prior experience with agentic AI systems. Figure 4.1 shows that most students had at least some experience using these tools, rather than being completely new users.

Participants were also asked whether they had ever paid to use an agentic AI system. As shown in Figure 4.2, the majority had reported they had not (63.2%, n=156). A smaller group was currently paying 27.1% (n=67), and 9.7% (n=24) reported that they had paid in the past.

Takeaway: These figures suggest while AI tools are widely used among this population, paid access is less—most students are working with free-tier tools.

Participants were also asked to compare agentic AI tools to more traditional sources of academic help. Figure 4.3 shows student perceptions of the extent to which an agentic AI could fill the roles of a professor, teaching assistant or tutor, and teammate on a group project. The pattern across all three roles is consistent—students drew a clear distinction between supplemental support and genuine replacement.

For the role of a professor, that distinction was the sharpest. The majority of participants reported that agentic AI could not fill this role to any extent (69.2%, n=200), 27% (n=78) indicated somewhat, and only 3.8% (n=11) reported completely.

For the role of a teaching assistant or tutor, responses were more evenly distributed. 52.6% (n=152) said somewhat, 27.7% (n=80) said not at all, and 19.7% (n=57) said completely. Suggesting students see AI as a more credible stand-in for peer or supplemental support than for formal instruction.

Lastly, the role of a teammate on a group project fell between these two extremes. Nearly half of participants reported agentic AI could not fill this role at all (49.5%, n=143), 33.9% (n=98) indicated somewhat, and 16.6% (n=48) said completely.

Takeaway: Most participants were already familiar with agentic AI tools, though paid access remained uncommon. Students were most willing to view AI as filling supplemental roles—tutoring especially—and far more skeptical of it replacing instructors entirely.

4.2 Usage Frequency

Participants reported how frequently they used agentic AI tools across both general and coursework-specific tasks. Figure 4.4 shows overall use frequency; Figure 4.5 breaks usage across task type.

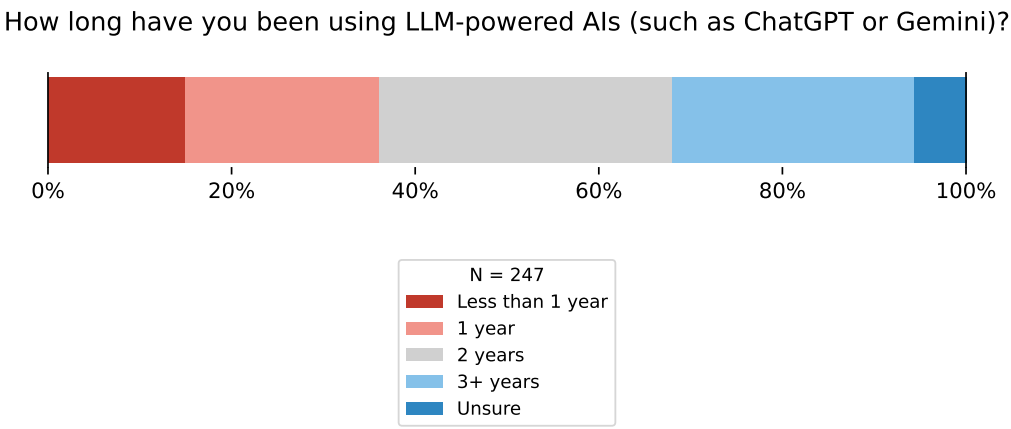


Figure 4.1: Length of participant experience using agentic AI tools.

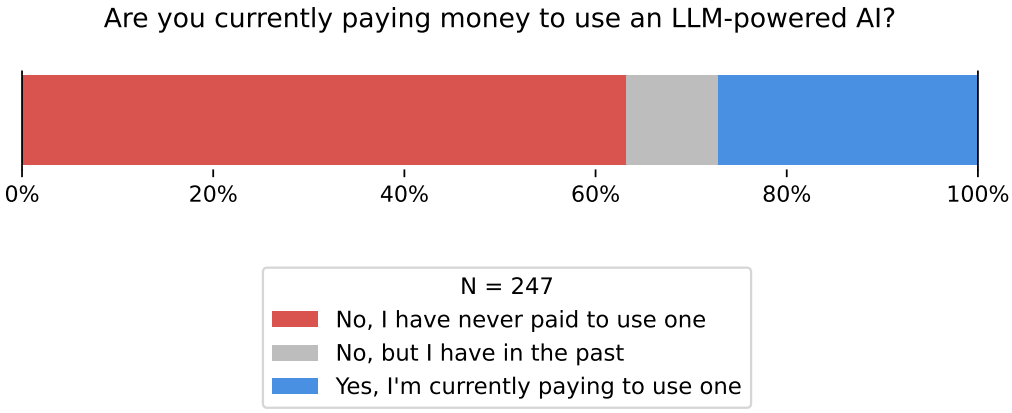


Figure 4.2: Reported paid use of agentic AI tools among participants.

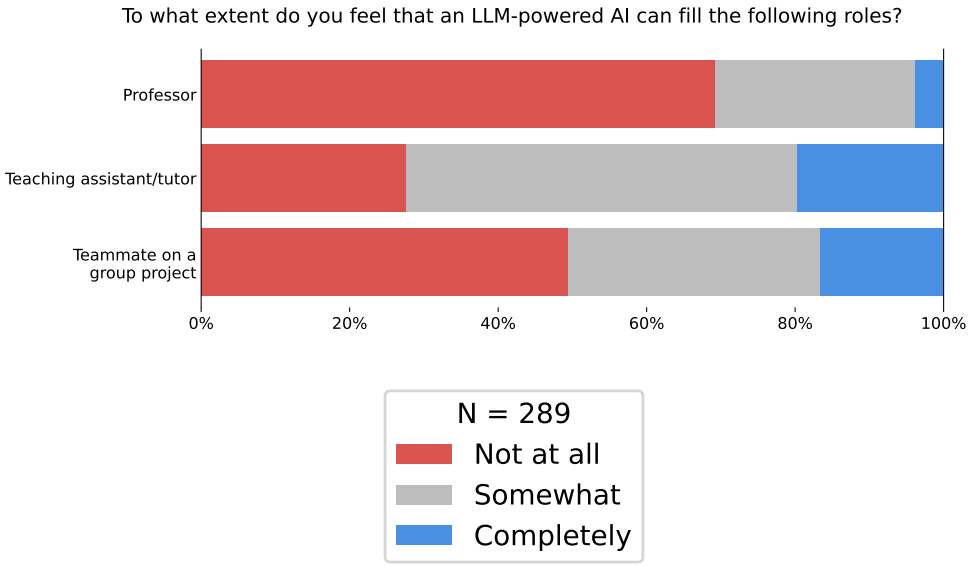


Figure 4.3: Student perceptions of the extent to which agentic AI can fill various academic roles.

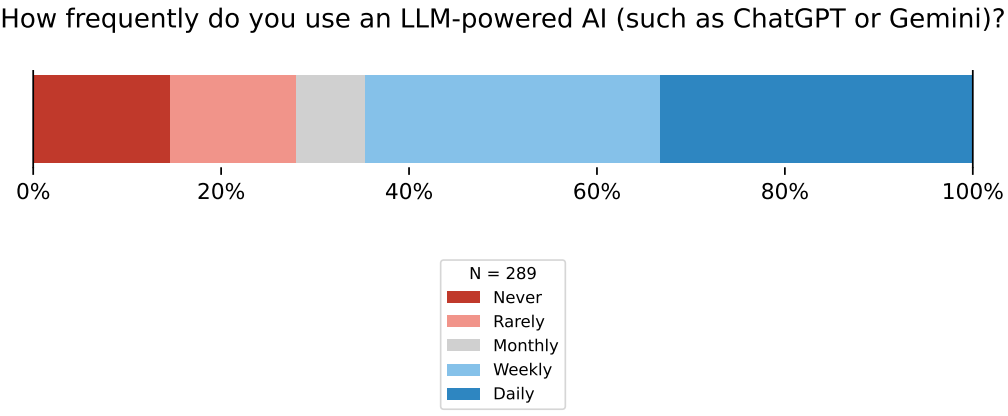


Figure 4.4: Reported frequency of agentic AI use across coursework-related tasks.

For general use, the most common responses were daily (33.2%, n=96) and weekly (31.5%, n=91), with smaller portions reporting never (14.5%, n=42), rarely (13.5%, n=39), or monthly (7.3%, n=21). These numbers suggest regular engagement—more than two-thirds of the sample is using AI tools at least weekly.

Task-specific patterns reveal where that engagement is concentrated. 4.5 shows the general use, the most common responses were debugging programming assignments, with 31.2% (n=77) reporting weekly use and 17.0% (n=42) reporting daily use. Understanding lectures and reading material follows a similar pattern, where responses were most distributed, with 29.1% (n=72) weekly and 13.8% (n=34) daily.

At the other end of the spectrum, non-code homework was the least supported task—40.5% (n=100) reported never using AI for it, and only 9.3% (n=23) reported daily use. Writing code for homework and exam preparation also trended lower, with 26.7% (n=66) and 30.8% (n=76) respectively reporting never. The pattern shows students reach for AI when they're stuck on something technical, less so when the work is conceptual or evaluative.

> **Takeaway:** These results suggest that agentic AI is used by students regularly. However, the usage frequency varies by task. Debugging and understanding course material saw the highest engagement, while non-code coursework and exam prep saw the least.

4.3 Perceived Usefulness

This section examines how students perceived the usefulness of agentic AI tools—both for general academic outcomes and specific-coursework related tasks. Figure 4.6 presents general usefulness perceptions, while Figure 4.7 focuses on task-specific helpfulness ratings.

Overall, students reported moderate to high perceived usefulness. The strongest positive responses came from two areas: completing work (40.9% (n=101) agreeing

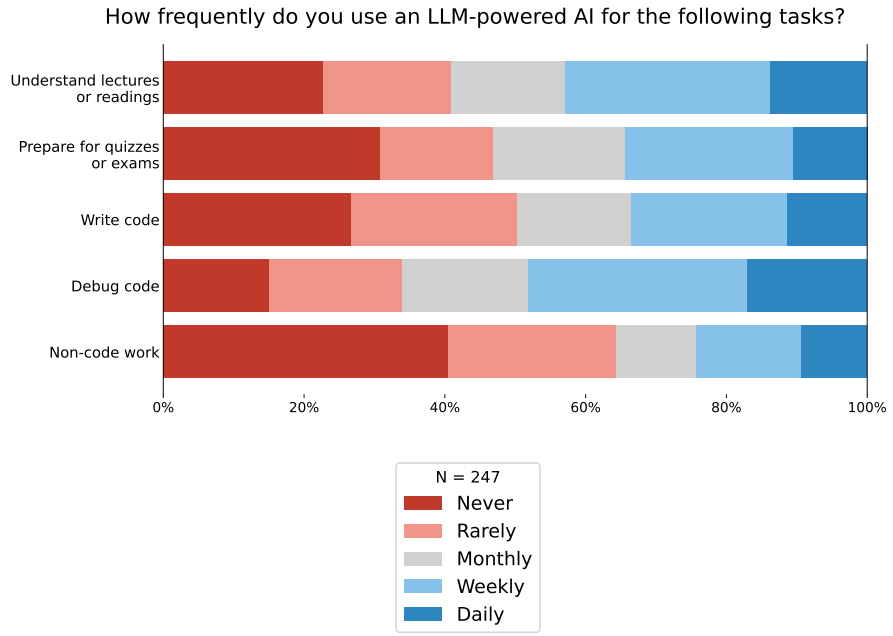


Figure 4.5: Reported frequency of agentic AI use across coursework-related tasks.

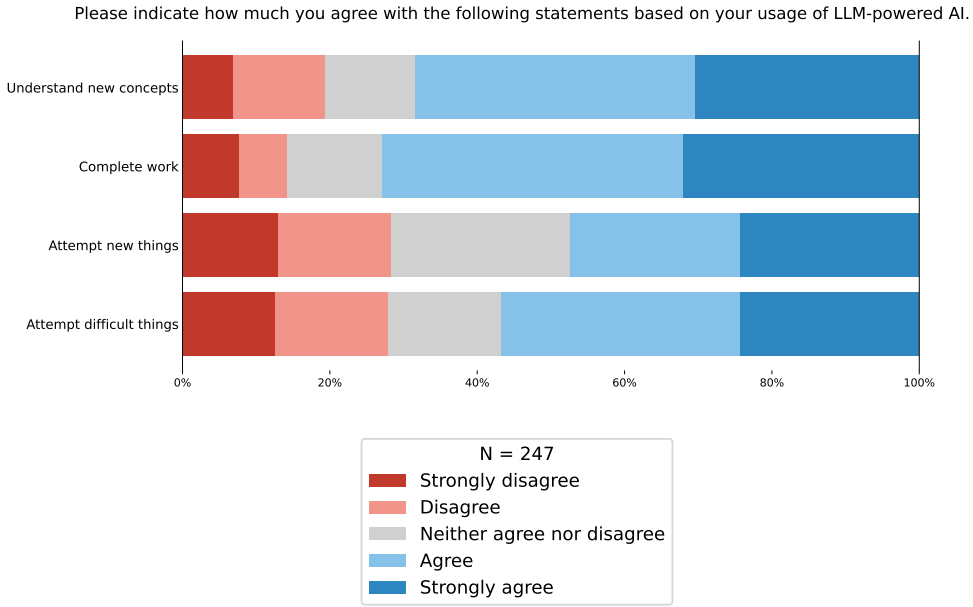


Figure 4.6: Student perceptions of the general usefulness of agentic AI tools.

and 32.0% (n=79) strongly agreeing) and understanding new concepts (38.1% (n=94) agreed and 30.4% (n=75) strongly agreed). Both suggest that students see AI as genuinely effective when the goal is getting something done or making sense of something unfamiliar.

Responses related to encouraging students to attempt more difficult tasks were slightly lower and more distributed across categories. 32.4% (n=80) of participants agreed and 24.3% (n=60) strongly agreeing, while 15.4% (n=38) were neutral and 28.0% (n=69) reported disagreement. For attempting new things, responses were more evenly distributed, with 24.3% (n=60) strongly agreeing, 23.1% (n=57) agreeing, and 24.3% (n=60) neutral.

Task-specific helpfulness followed the same pattern as usage frequency. Debugging was rated the most helpful application, with 34.3% (n=72) reporting extremely helpful and 27.6% (n=58) very helpful. Understanding lectures and material was also rated highly, with 27.6% (n=53) reporting extremely helpful and 25.1% (n=48) very helpful.

Other tasks show slightly lower but still positive responses. For writing code, responses were more distributed, with 26.5% (n=48) reporting moderately helpful and similar proportions across other categories. For quizzes and exams, 28.1% (n=48) reported very helpful and 26.3% (n=45) extremely helpful. Non-code work showed the most balanced distribution, with responses spread across somewhat helpful (26.5%, n=39), moderately helpful (24.5%, n=36), very helpful (22.4%, n=33), and extremely helpful (22.4%, n=33).

These findings align closely with the usage frequency data from the previous section. Where students use AI most, they also find it most useful—which reinforces the picture of AI as practical, task driven tool rather than a general purpose academic resource.

Takeaway: Students generally perceived agentic AI as useful, particularly when completing work, understanding new concepts, and debugging code. Perceived usefulness was lower—and more divided for tasks that required risk-taking or working outside technical domains.

4.4 Trust and Reliability

Participants were asked how confident they were in the accuracy and reliability of AI-generated outputs. This measure is shown as part of the broader set of agreement-based statements in Figure4.8, specifically under *"Confidence in answer accuracy"* item.

The most common responses were strongly disagree (27.5%, n=68) and neither agree nor disagree (27.5%, n=68), followed by disagree (20.6%, n=51). Only 17.8% (n=44) agreed and 6.5% (n=16) strongly agreeing that AI-generated answers are accurate. Compared to most other statements, in Figure 4.8, the accuracy item is notably more spread across the scale.

This contrasts with the higher levels of perceived usefulness reported in the previous section. Students generally view agentic AI as helpful for completing work and learning new material, but they do appear to be more cautious when evaluating whether the information provided is correct [12, 2]. The implication is that students are using AI as a scaffold rather than as a source to defer to. That's more than a simple "use it or don't" framing would suggest.

Takeaway: Despite high perceived usefulness, students expressed low to moderate confidence in the accuracy of AI outputs. The gap between helpfulness and trust suggests students are using AI as a support tool while maintaining independent judgment about whether its answers are correct.

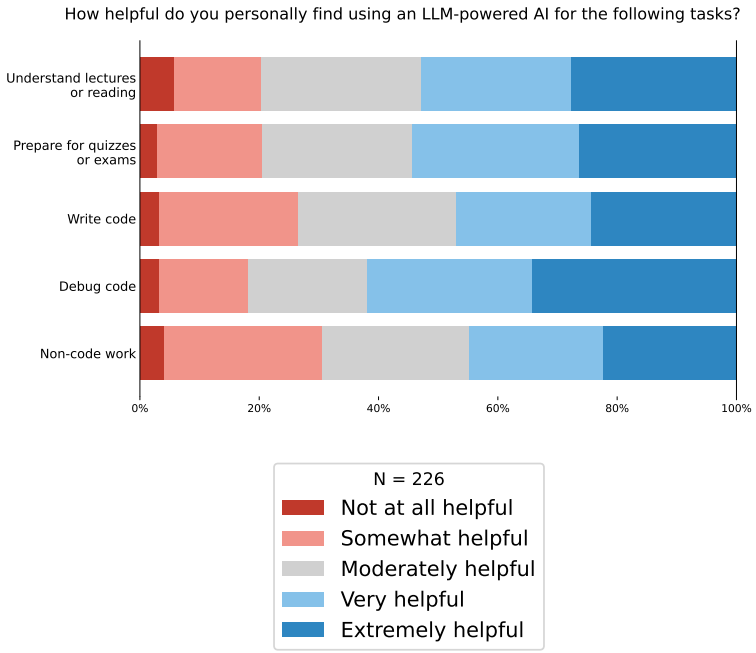


Figure 4.7: Student perceptions of the usefulness of agentic AI tools across coursework-related tasks.

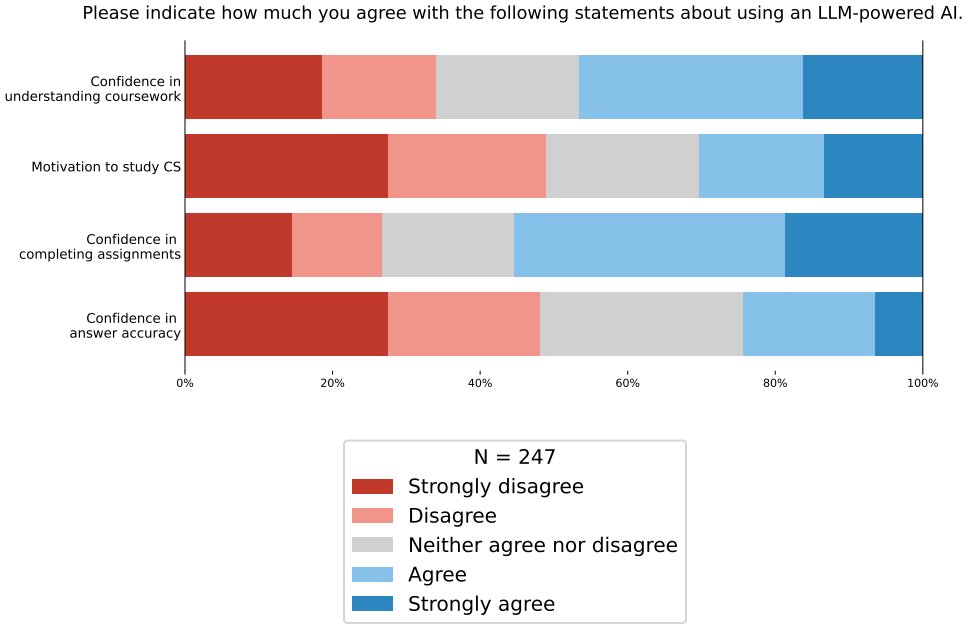


Figure 4.8: Student confidence in the accuracy and reliability of agentic AI outputs.

4.5 Impact on Learning and Motivation

Participants were asked about the impact of agentic AI on their learning experience—specifically changes to confidence, motivation, and their ability to complete coursework. These responses are also shown in Figure 4.8.

The impact on confidence was broadly positive. For understanding coursework, 30.4% (n=75) agreed and 16.2% (n=40) strongly agreed that agentic AI increased their confidence in understanding course material. For completing assignments, the numbers were slightly stronger: 36.8% (n=91) agreed and 18.6% (n=46) strongly agreed. Students appear to feel more capable —both of understanding and finishing their work—when using agentic AI tools.

The motivational picture is different. For motivation to study computer science, the most common responses were strongly disagree (27.5%, n=68) and disagree (21.5%, n=53), with 20.6% (n=51) selected neutral. Only 17.0% (n=42) agreed and (13.4%, n=33) strongly agreed. AI, it seems, helps students get through the work, but it doesn't make them want to do it more.

This distinction matters. Confidence and task completion are proximate outcomes—they show up in the moment of using the tool. Motivation operates on a longer timescale, shaping how students engage with the field overall. The fact AI appears to support the former without meaningfully affecting the latter suggests a gap between what these tools do well and what deeper learning requires.

> **Takeaway:** Agentic AI has a clearer positive effect on confidence and task completion than on motivation. Students felt more capable when using AI tools, but were less likely to report feeling more motivated to engage with their coursework.

4.6 Perceptions of AI in Future Careers

Participants were asked about the role they expected AI to play in their future careers in computer science. These responses are shown in Figure 4.9.

On the question of job difficulty, responses leaned clearly in one direction: 41.2% (n=119) agreed and 29.4% (n=85) strongly agreed that AI will make it harder to obtain a job in computer science. 14.5% (n=42) were neutral, while 11.4% (n=33) disagreed and 3.5% (n=10) strongly disagreed.

On whether future computer science jobs will require AI tools, the distribution is more balanced: 34.9% (n=101) agreed and 20.1% (n=58) strongly agreed, but 18.0% (n=52) selected neither agree nor disagree, 17.3% (n=50) disagreed, and 9.7% (n=28) strongly disagreed. There is a rough consensus that AI proficiency will matter—but not a unanimous one.

The most striking result was on enthusiasm. When asked whether they were looking forward to using AI in future careers, the two most common responses were strongly disagree (30.1%, n=87) and disagree (22.5%, n=65), with (18.0%, n=52) neutral. Only (14.9%, n=43) agreed, and (14.5, n=42) strongly agreed. For a population that uses AI regularly and finds it useful, the practical value of these tools does not appear to translate into enthusiasm about a career shaped by them.

> **Takeaway:** Students expressed mixed views about AI's role in their futures. Most expected AI to become increasingly important in the workforce, and many were concerned it would make the job market more competitive. Despite that, enthusiasm about working with AI long-term was limited.

4.7 Qualitative Findings

Open-ended responses provided texture to the quantitative results—not just confirming the patterns, but explaining the reasoning behind them. There were several recurring themes across participants.

Useful, but needs limits. The most consistent theme was that AI is genuinely a valuable tool, but should not replace independent thinking. Many students emphasized that the value of these systems depends heavily on how they are used.

For example, one participant stated,

> *"If used correctly, it can be a great tool, but users still need to be great without it."(P2)*

Similarly, another noted that AI should be used under supervision,

> *"We should not let models take the wheel. We should use them to enrich our reasoning rather than to replace our reasoning."(P26)*

This perspective aligns with quantitative findings, where students reported high perceived usefulness, but lower levels of trust and confidence in AI-generated outputs.

Critical thinking concerns. Many students expressed concern that reliance on AI may negatively impact critical thinking and long-term learning. Several responses described AI as a tool that can *"severely affect your ability to think critically and can destroy any chance you have at learning from new problems or topics that you have never encountered before"(P138)*, or foster a dependency that undermines long-term learning. This reinforces the motivational gap observed in Section 4.5—AI may be helping students complete work without necessarily deepening their engagement with it.

Trust and correctness issues. Accuracy concerns that appeared in the Likert data showed up in the open-ended responses too, often with sharper language. One participant described AI outputs as, *"by their nature they are often confidently wrong"(P130).* Another called it *"a jack of all trades rather than a skilled specialist,"(P25)* particularly in more advanced contexts. These responses give concrete texture to the low-trust numbers from Section 4.4.

Ethics and societal concerns. A significant portion of responses focused on ethical and societal concerns surrounding AI. Participants raised concerns related to training data, environmental impact, misinformation, and societal consequences of widespread AI adoption. One participant described AI systems as *"fundamentally built on theft,"(P104)*; another elaborated on concerns about,

> *"How it relies on mass copyright infringement for training data, the massive amount of power and resources it consumes, and so forth."(P131)*

Others highlighted the potential for misinformation, reduced societal trust, and negative impacts on education. Students are thinking at this level—while using these tools daily—points to a genuine tension between practical adoption and ethical unease.

Career uncertainty. Career-related responses reflected the same split observed in the quantitative data: a mix of optimism and concern. Some students saw AI proficiency as a competitive advantage; others were anxious about what it meant for employment. The range is captured well in two responses side by side:

> *"I personally don't think AI will take people's jobs but the people not using it/embracing it may see a loss of work."(P82)*

> *"Please keep hiring recent CS grads :,-)"(P12)*

The humor in P12's response shouldn't obscure what it's actually expressing—a real anxiety about whether the skills students are building will still be valued by the time they graduate.

> **Takeaway:** Qualitative responses reinforced the quantitative findings while also adding nuance the numbers alone couldn't capture. Students described AI as useful but limited, raising recurring concerns about trust, critical thinking, ethics, career uncertainty, and emotional discomfort.

4.8 Summary of Results

Overall, the results show that students are already integrating agentic AI tools into their academic workflows—particularly for technical work such as debugging and understanding course material—and they broadly find those tools useful. More than two-thirds of participants reported using AI at least weekly, and perceived usefulness was high for task completion and conceptual understanding.

That usefulness, however, does not extend to trust. Students reported relatively low confidence in the accuracy of AI outputs, with many expressing uncertainty or disagreement when asked whether agentic AI outputs are reliable—a finding across both the Likert data and the open-ended responses. The clearest way to describe the relationship students have with these tools: *helpful, but verify.*

The learning impact data adds another layer. AI appears to have a stronger influence on confidence and task completion than on motivation. While many students reported feeling more capable of understanding material and completing assignments, fewer indicated that AI increased their motivation to engage with coursework.

Finally, perceptions of AI's role in future careers reflect a mix of recognition and concern. Students generally believe that AI will become an important part of computer science jobs, but many also feel that it may make obtaining those jobs more difficult.

Additionally, enthusiasm about using AI in future careers was more limited, suggesting uncertainty about its long-term role.

These findings set up the discussion chapter directly: the question is no longer whether students are using AI, but what that use actually means for learning, trust, motivation, and the shape of a CS education.

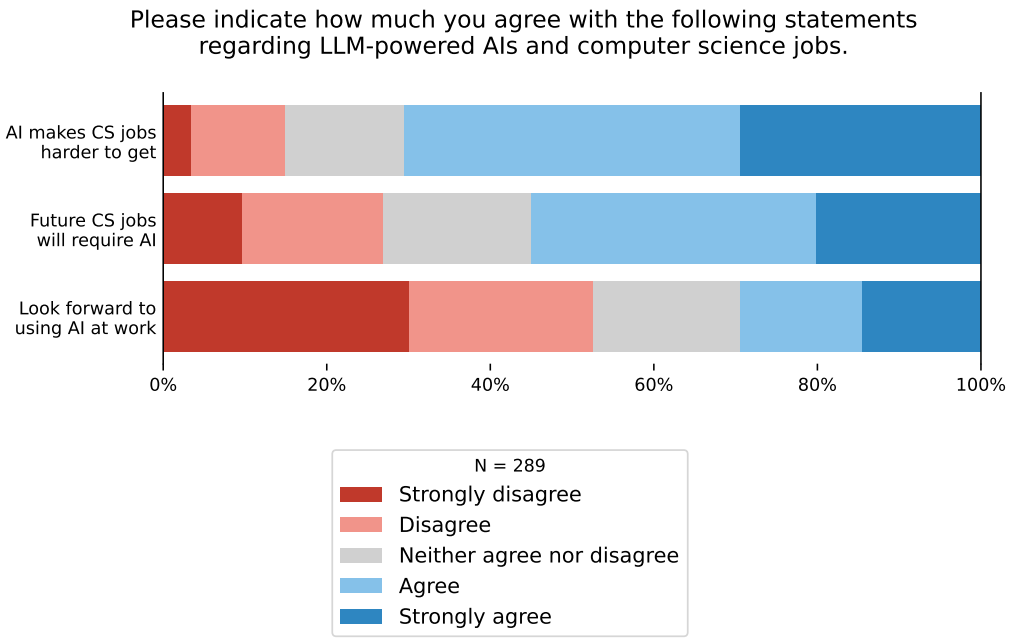


Figure 4.9: Student perceptions of the role of agentic AI tools in future computer science careers.

Chapter 5

Discussion

Students are using agentic AI tools—regularly, purposefully, and with a fair amount of nuance about what those tools can and cannot do.

5.1 What the Data Tells Us

Students are integrating agentic AI into their academic workflows at a notable rate, with usage concentrated around technical tasks like debugging and understanding course material. Engagement is moderate but consistent—more than two-thirds of the participants reported using these tools at least weekly, depending on the task.

Perceived usefulness is high, particularly for completing work and understanding new concepts. That is not a surprising finding on its own. What makes it interesting is what sits alongside it: low to moderate confidence in whether AI-generated outputs are actually correct. Students find these tools helpful and don't fully trust them—simultaneously.

On learning impact, AI appears to support confidence and task completion; it does not appear to move the needle on motivation. Students feel more capable when using these tools, but not more engaged. On future careers, there is a tension that runs through the data consistently, students expect AI to become important in the

workforce, worry it will make hiring more competitive, and—despite all of that—are not particularly enthusiastic about a career defined by it.

One additional note on the statistical analysis—tests of association were conducted to examine relationships between demographic variables and survey responses. Multiple comparison corrections were applied using the Benjamini-Yekutieli method to control a false discovery rate. No relationships remained statistically significant after correction. This was an intentionally conservative approach—the goal was to avoid surfacing patterns that might be noise. The absence of significant findings is reported here transparently, rather than omitted.

5.2 The Trust Gap

The most theoretically interesting finding in this study is the gap between usefulness and trust—and the fact that it doesn't close with experience. Students who had been using agentic AI tools longer did not report higher confidence in output accuracy.

A reasonable prior assumption might be that students who use these tools more would calibrate their trust upward over time—learning which prompts work, which outputs to verify, and when to rely on AI versus their own judgment. The data suggests something different, however. The increased exposure may actually reinforce awareness of AI's limitations rather than reduce it.

Qualitative responses reinforce this interpretation, with many students describing AI as *"confidently wrong"(P130, P145, P281)* or emphasizing the need to verify outputs independently. That experience, repeated over time, appears to produce a stable mental model. *AI is useful, but is inherently unreliable.* That is just the property of the tool. One participant stated AI responses should *"never be trusted or taken at face value,"(P130)* highlighting a strong awareness of potential inaccuracies.

Occam's razor suggests that the trust gap may not be a problem to solve. Educational environments might do better helping students develop verification habits than trying to build trust in outputs directly.

5.3 Implications for Learning

The confidence versus motivation split deserves more attention than it might initially appear to warrant. Confidence and task completion show up in the moment of using the tool. Motivation operates on a longer timescale—shaping whether students seek out harder problems, engage more deeply with the material, and sustain interest in the discipline. The fact AI supports the former without affecting the latter suggests a gap between what these tools do well and what deeper learning actually requires.

The qualitative data sharpens this concern. Students weren't just worried about over-reliance— several described experiencing it directly:

> *"I think that using LLM-powered AI is making me less intelligent ... I feel like I am too reliant on AI to write for me."(P111)*

> *"I can feel myself becoming a worse coder and problem solver the more I use LLMs. It takes great self restraint not to use, especially within CS, when simple debugging help can quickly morph into entire projects coded predominantly by AI."(P164)*

These responses suggest that students are not only aware of potential learning trade offs, but are actively experiencing them. That is a different kind of concern than general anxiety about AI in education. It points to a tension between efficiency and productive struggle—where the ease of getting to an answer reduces the cognitive work that builds lasting skill. In computer science especially, where debugging and problem decomposition are themselves core elements of the education—that trade-off is particularly worth taking seriously.

5.4 The Emotional Dimension

Something in the qualitative data that does not appear in the numbers is a recurring emotional register that accompanied AI use. Not just skepticism or rationality, but shame, guilt, anxiety, and dependency. Feelings that suggest *AI use has become entangled with students' sense of academic identity.*

> *"I actively feel shame every time I use it,"(P44)*

> *"I feel like I am becoming too dependent on it and I can't help being surrounded by it, I'm worried."(P222)*

These responses raise questions that quantitative measures cannot fully address. If students are using AI tools regularly while also feeling guilty about it, that emotional friction is doing something. It is shaping how they evaluate the outputs, how honestly they report their usage, and how they think about their own competence. A student who feels shame about using AI may also under report it on a survey, which is worth noting as a limitation of self-reported data.

Educational strategies that address only the functional use of AI—when it's permitted, how to cite it, what counts as academic dishonesty—may be missing a layer that actually drives the behavior.

5.5 Limitations

Several limitations should be acknowledged before drawing a set of conclusions from these findings.

The study relies entirely on self-reported data. Participants may over or underestimate their AI usage. The emotional dynamics described in section 5.4 suggest that under reporting in particular is plausible. This concern is compounded by the recruitment method itself.

Participants who attend institutions with strict AI policies might have been biased to under report their AI usage—even though no data was collected. The survey invitation arrived through official school channels, which may have discouraged honest participant answers—regardless of whether data was collected or not. What students say they do and what they actually do with these tools may not fully align.

The sample, while broad in institutional coverage, skews heavily toward university students in the United States—and within that, towards higher-GPA students. These findings may not generalize to community college populations, students outside the U.S., or students in non-computer science disciplines, where the student-AI relationship may look quite different.

The study also captures a single point in time. Given how rapidly agentic AI tools are evolving, student attitudes may shift quickly. The trust gap observed here might narrow or widen as these tools improve and the norms around them shift.

Finally, the coding of qualitative responses introduces a degree of subjectivity. Efforts were made to apply Boolean indicators consistently, however the interpretation of open-ended responses is never fully objective.

5.6 Recommendations

For educators. The findings suggest that students are already using agentic AI tools regularly, regardless of formal guidance. As a result, educators should focus on teaching students how to use these tools effectively and the consequences of overuse, rather than attempting to restrict their use entirely.

Clear guidelines should be established to distinguish between acceptable assistance and academic dishonesty, as emphasized in prior work on AI and academic integrity [6, 23]. Additionally, assignments may need to be adapted to emphasize understanding and reasoning rather than final outputs alone.

For students. Students should be encouraged to use agentic AI as a support tool rather than a replacement for learning. This includes using AI for debugging, clarification, and exploration—while still engaging in independent problem solving.

Developing the ability to critically evaluate AI-generated outputs is also essential, particularly given the known risks of hallucinated or incorrect responses [12].

For researchers. Future research should investigate the long-term effects of AI usage on learning outcomes, including retention, problem-solving ability, and academic performance. Longitudinal studies may provide insight into how these effects evolve over time. Additionally, further work is needed to better understand the relationship between trust, usage patterns, and learning behaviors.

For AI companies. Companies should consider improving transparency and reliability in AI systems, particularly in educational contexts. Features that help users verify information or understand the reasoning behind outputs may help address the trust gap identified in this study [3].

For institutions. Institutions should develop clear policies regarding AI usage that balances academic integrity with practical adoption. Policies should be communicated clearly to students and aligned with evolving technological capabilities.

5.7 Future Research Directions

The findings of this study highlight several opportunities for future research aimed at better understanding the role of agentic AI in computer science education.

First, longitudinal studies are needed to examine how student perceptions and usage patterns evolve over time. As agentic AI tools continue to improve and become more deeply integrated into academic and professional environments, it is important to determine whether attitudes toward trust, usefulness, and dependency shift with prolonged exposure. In particular, the persistent trust gap observed in this study

raises questions about whether trust can be developed over time [10], or whether skepticism persists even with increased exposure, as suggested in prior work on AI reliability and user trust [2, 12].

Secondly, expanding this work to include a broader range of disciplines and geographic regions would improve the generalization of these findings. While this study focuses on computer science, students in the United States; students in other fields may interact with agentic AI in fundamentally different ways. Particularly in disciplines that rely less heavily on technical problem solving.

Additionally, future research should investigate the relationship between AI usage and measurable learning outcomes. While this study focuses on student perceptions, it remains unclear whether the use of agentic AI tools supports or hinders long-term learning, including retention, conceptual understanding, and independent problem-solving abilities.

Further qualitative research, such as interviews or focus groups, could provide deeper insight into how students perceive ethical concerns, issues of trust, and dependency. The strong emotional responses observed in this study suggest that student interactions with AI are not purely functional, but are shaped by broader psychological and social factors that warrant closer examination.

Finally, future work may explore how institutional policies, instructor methods, and classroom guidelines influence student behavior and perceptions regarding the use of agentic AI. Understanding how different approaches to AI integration affect learning outcomes and student experiences will be critical for developing effective educational strategies.

Chapter 6

Conclusions

This study examined how computer science students perceive and use agentic AI tools in their academic coursework. The answer: regularly, practically, and with considerably more nuance than either enthusiastic adoption or blanket skepticism would suggest.

Overall, students reported moderate to high levels of perceived usefulness, indicating that agentic AI tools are widely viewed as effective for supporting productivity and task completion. However, this usefulness does not fully translate into trust. Students expressed lower confidence in the accuracy and reliability of AI-generated outputs, highlighting a consistent gap between using AI as a tool and reliability of responses.

This trust gap does not appear to diminish with increased experience. Students who reported longer-term use of agentic AI tools did not demonstrate higher confidence in their reliability. This suggests that the adoption of these tools is driven more by convenience and efficiency rather than by trust in their outputs.

In terms of learning, agentic AI appears to have a stronger impact on confidence and task completion than on motivation. While many students reported feeling more capable of completing assignments, fewer indicated that these tools increased their motivation to engage with coursework. This highlights the tension between efficiency and deeper learning.

Qualitative responses further revealed that student interactions with AI are not purely functional. Many participants expressed concerns about over-reliance, critical thinking, and long-term skill development, while others described emotional responses such as guilt, discomfort, and dependency. These findings suggest that the impact of agentic AI extends beyond performance, influencing how students think about their learning and their role within it.

Perceptions of AI in future careers were similarly mixed. Students recognized that AI will likely play an important role in the workforce, but also expressed concerns about their potential impact on job opportunities and skill development.

Taken together, these results suggest that agentic AI tools are not replacing traditional learning, but are instead reshaping how students engage with it. While these systems provide clear benefits in terms of efficiency and accessibility, their long-term impact on learning remains dependent on how they are integrated and used.

As agentic AI continues to evolve, understanding student perceptions will be critical for guiding its role in computer science education. Future work should continue to explore how these tools influence learning outcomes over time, and how educators and institutions can support their responsible and effective use.

Bibliography

[1] B. A. Becker, P. Denny, J. Finnie-Ansley, A. Luxton-Reilly, J. Prather, and E. A. Santos. Programming is hard - or at least it used to be: Educational opportunities and challenges of ai code generation. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*, SIGCSE 2023, page 500–506, New York, NY, USA, 2023. Association for Computing Machinery. 1

[2] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA, 2021. Association for Computing Machinery. 6, 21, 36

[3] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill, E. Brynjolfsson, S. Buch, D. Card, R. Castellon, N. S. Chatterji, A. S. Chen, K. A. Creel, J. Davis, D. Demszky, C. Donahue, M. Doumbouya, E. Durmus, S. Ermon, J. Etchemendy, K. Ethayarajh, L. Fei-Fei, C. Finn, T. Gale, L. E. Gillespie, K. Goel, N. D. Goodman, S. Grossman, N. Guha, T. Hashimoto, P. Henderson, J. Hewitt, D. E. Ho, J. Hong, K. Hsu, J. Huang, T. F. Icard, S. Jain, D. Jurafsky, P. Kalluri, S. Karamcheti, G. Keeling, F. Khani, O. Khattab, P. W. Koh, M. S. Krass, R. Krishna, R. Kuditipudi, A. Kumar, F. Ladhak, M. Lee, T. Lee, J. Leskovec, I. Levent, X. L. Li, X. Li, T. Ma, A. Malik, C. D. Manning, S. P. Mirchandani, E. Mitchell, Z. Munyikwa, S. Nair, A. Narayan, D. Narayanan, B. Newman, A. Nie,

J. C. Niebles, H. Nilforoshan, J. F. Nyarko, G. Ogut, L. Orr, I. Papadimitriou, J. S. Park, C. Piech, E. Portelance, C. Potts, A. Raghunathan, R. Reich, H. Ren, F. Rong, Y. H. Roohani, C. Ruiz, J. Ryan, C. R'e, D. Sadigh, S. Sagawa, K. Santhanam, A. Shih, K. P. Srinivasan, A. Tamkin, R. Taori, A. W. Thomas, F. Tramèr, R. E. Wang, W. Wang, B. Wu, J. Wu, Y. Wu, S. M. Xie, M. Yasunaga, J. You, M. A. Zaharia, M. Zhang, T. Zhang, X. Zhang, Y. Zhang, L. Zheng, K. Zhou, and P. Liang. On the opportunities and risks of foundation models. *ArXiv*, 2021. 6, 35

[4] N. Carr. *The Shallows: What the Internet Is Doing to Our Brains*. W. W. Norton, 2010. 5

[5] College Board. Bigfuture college search, 2024. Last accessed: 2026-01. 8, 10

[6] D. R. E. Cotton, P. A. Cotton, and J. R. Shipway. Chatting and cheating: Ensuring academic integrity in the era of chatgpt. *Innovations in Education and Teaching International*, 61(2):228–239, 2024. 5, 34

[7] L. e. a. Floridi. Ai4people—an ethical framework for a good ai society. *Minds and Machines*, 28:689–707, 2018. 6

[8] M. Gerlich. Ai tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 2025. 2, 5, 6

[9] S. Giannini. Generation ai: Navigating the opportunities and risks of artificial intelligence in education, 2024. UNESCO, keynote-based article, accessed 2026. 1, 3

[10] E. Glikson and A. W. Woolley. Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2):627–660, 2020. 36

[11] W. Holmes, M. Bialik, and C. Fadel. *Artificial Intelligence in Education. Promise and Implications for Teaching and Learning.* Center for Curriculum Redesign, 03 2019. 2, 3, 4

[12] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. J. Bang, A. Madotto, and P. Fung. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12), Mar. 2023. 6, 21, 35, 36

[13] C. K. Lo. What is the impact of chatgpt on education? a rapid review of the literature. *Education Sciences*, 13(4), 2023. 4

[14] D. Y. Long, S. Wang, S. Md Rashid, and X. T. Lu. Artificial intelligence in higher education: a systematic review of its impact on student engagement and the mediating role of teaching methods. *Frontiers in Education*, Volume 10 - 2025, 2026. 2, 3

[15] A. Machado, A. Pesqueira, and M. Sousa. *A Review ChatGPT and the Future of Education*, pages 155–176. Springer, 03 2024. 2

[16] B. N. Macnamara, I. Berber, M. C. Çavuşoğlu, E. A. Krupinski, N. Nallapareddy, N. E. Nelson, P. J. Smith, A. L. Wilson-Delfosse, and S. Ray. Does using artificial intelligence assistance accelerate skill decay and hinder skill development without performers' awareness? *Cognitive Research: Principles and Implications*, 9(1):46, 2024. 2, 5, 6

[17] W. K. Monib, A. Qazi, R. A. Apong, M. T. Azizan, L. De Silva, and H. Yassin. Generative ai and future education: a review, theoretical validation, and authors' perspective on challenges and solutions. *PeerJ Computer Science*, 10:e2105, 2024. 1, 4, 5

[18] Northwestern University School of Education and Social Policy. Closing the gender gap in computer science, March 2024. Accessed: 2026-04-21. 11

[19] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 6

[20] I. C. Pelaez-Sanchez, D. Velarde-Camaqui, and L. D. Glasserman-Morales. The impact of large language models on higher education: exploring the connection between ai and education 4.0. *Frontiers in Education*, 9, 2024. 1, 4, 6

[21] D. Ravšelj, D. Keržič, N. Tomaževič, L. Umek, N. Brezovar, A. Aristovnik, et al. Higher education students' perceptions of chatgpt: A global study of early reactions. *PLOS ONE*, 20(2):e0315011, 2025. 1, 2, 4, 5, 6

[22] B. Sparrow, J. Liu, and D. M. Wegner. Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043):776–778, 2011. 5, 6

[23] T. Susnjak and T. R. McIntosh. Chatgpt: The end of online exam integrity? *Education Sciences*, 14(6), 2024. 34

[24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 1

[25] X. Zhai. Chatgpt for next generation science learning. *XRDS*, 29(3):42–46, Apr. 2023. 4, 5

Vita

Margaret Kelley was born and raised in Sevierville and Knoxville, Tennessee, alongside her brother Charles. Her parents, Sharon and John Kelley, supported her in sports such as karate and horseback riding, and in creative pursuits through art classes and supplying endless amounts of yarn for knitting projects.

She was homeschooled until her Junior year of high school, where she attended and graduated The L&N STEM Academy. This is where she took her first computer science class, with Mr. Jason Colby, who inspired her to continue in computer science.

She then went on to graduate with an A.S. in Science from Pellissippi Community College in 2022.

She graduated with her B.S. in Computer Science and minor in Cybersecurity from The University of Tennessee.

In her undergraduate degree, she began research with the USER Lab, overseen by Dr. Scott Ruoti.

This encouraged her to pursue her life long dream of graduate school so she could become a teacher.