

Evaluating CS1-LLM: Integrating LLMs and Examining Student Outcomes in an Introductory Computer Science Course

Annapurna Vadaparty
UC San Diego
La Jolla, CA, USA
avadaparty@ucsd.edu

David H. Smith IV
Virginia Tech
Blacksburg, VA, USA
dhsmith4@vt.edu

Samvrit Srinath
UC San Diego
La Jolla, CA, USA
sasrinath@ucsd.edu

Mounika Padala
UC San Diego
La Jolla, CA, USA
mpadala@ucsd.edu

Christine Alvarado
UC San Diego
La Jolla, CA, USA
cjalvarado@eng.ucsd.edu

Jamie Gorson Benario
Education for Social Impact
Google
Chicago, IL, IL, USA
jamben@google.com

Leo Porter
UC San Diego
La Jolla, CA, USA
leporter@eng.ucsd.edu

Daniel Zingaro
University of Toronto Mississauga
Mississauga, Ontario, Canada
daniel.zingaro@utoronto.ca

Abstract

Large language models (LLMs) have broad implications for education in general, impacting the foundations of what we teach and how we assess. This is especially true in computing, where LLMs tuned for coding have demonstrated shockingly good performance on the types of assignments historically used in introductory CS (CS1) courses. As a result, CS1 courses will need to change in terms of the skills that are taught and how they are assessed. Computing education researchers have begun to study student use of LLMs, but there remains much to be understood about the ways that these tools affect student outcomes. In this paper, we present the design and evaluation of a new CS1 course at a large research-intensive university that integrates the use of LLMs for student learning. We describe the design principles used to create our course, our new course objectives, and evaluation of student outcomes and perceptions throughout the course as measured by assessment scores and surveys. Our findings suggest that 1) student exam performance outcomes, including differences among demographic groups, are largely similar to historical outcomes for courses without integration of LLM tools, 2) large, open-ended projects may be particularly valuable in an LLM context, and 3) students predominantly found the LLM tools helpful, although some had concerns regarding over-reliance on the tools.

CCS Concepts

• **Social and professional topics** → **Computing education**;
Software engineering education; CS1.

Keywords

Introductory Programming, CS1, GenAI, LLM, Student Outcomes

ACM Reference Format:

Annapurna Vadaparty, David H. Smith IV, Samvrit Srinath, Mounika Padala, Christine Alvarado, Jamie Gorson Benario, Leo Porter, and Daniel Zingaro. 2026. Evaluating CS1-LLM: Integrating LLMs and Examining Student Outcomes in an Introductory Computer Science Course. In *28th Australasian Computing Education Conference (ACE 2026), February 09–13, 2026, Melbourne, VIC, Australia*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3786228.3786234>

1 Introduction

If the advent of the Internet beckoned the information age, the advent of widely available large language models (LLMs) represents our entering the age of AI. The widespread availability of these models has led to concerns over academic integrity [16, 31]. In computing education, these concerns are made more acute by the widespread availability of LLM-powered programming assistants.

The role of LLMs in computing courses is under intense study [41], but there has been little research on the outcomes of fully redesigned courses integrating LLMs. Current studies have focused on small-scale interventions, student perceptions, and behavior when using LLM tools [21, 56]. While these studies provide preliminary insights, the impact of LLM integration in large-scale university courses remains unclear.

We have designed a new introductory CS (CS1) course that fully incorporates student use of LLM tools. In this paper, we investigate how student outcomes in this course compare to historical norms, both with respect to international benchmarks [45] and across demographic groups that have been of historical interest to computing education [17, 19, 34, 42, 51]. To guide future iterations of courses that may seek to adopt our approach, we evaluate student perceptions of the course and the LLM tools they employed. To these ends, we address the following research questions:

RQ1: How did students' performance when learning to program with LLM tools compare to historical, pre-LLM benchmarks?



This work is licensed under a Creative Commons Attribution 4.0 International License.
ACE 2026, Melbourne, VIC, Australia
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2352-0/26/02
<https://doi.org/10.1145/3786228.3786234>

RQ2: What were students’ perceptions of learning to program with LLM tools?

RQ3: How does student performance vary across different student populations in a course that centers LLMs?

2 Literature Review

2.1 Large Language Models in CS1

The advent of widely available Large Language Models (LLMs) is reshaping introductory computer science education. If an LLM were a student, it would rank among the top performers in CS1 assessments [15]. Furthermore, the widespread adoption of tools like GitHub Copilot in industry, used by over 20 million developers¹, signals that students entering the workforce will need proficiency in these tools. As a result, traditional educational approaches are being reexamined, as discussed by Becker et al. [4] in “*Programming is Hard – Or at Least it Used to Be.*”

“AI-generated code presents both opportunities and challenges for students and educators in introductory programming and related courses. The sudden viability and ease of access to these tools suggest educators may be caught unaware or unprepared for the significant impact on education practice resulting from AI-generated code. We therefore urgently need to review our educational practices in the light of these new technologies.”

They, and others, suggest that the rise of LLMs will give way to novel pedagogical approaches that reduce barriers to programming and allow students to focus on higher-level concepts such as problem decomposition, algorithmic design, and developing larger applications earlier in their education [13, 39]. While not a panacea for all the barriers students face when learning to program, these models, and systems built on top of them, have the opportunity to offer additional support to students through AI tutoring systems [21, 26], auto-grading systems [1, 49], and content generation [27].

Prior studies have explored students’ use of LLM tools in computing, such as through prompt-writing tasks that require understanding a programming problem to prompt functional code [11, 12, 48]. Lab-based (quasi-experimental) studies have observed student behaviors and surveyed attitudes toward using LLM tools in programming tasks [21, 56], showing generally positive attitudes and similar performance compared to peers without LLM tools. However, these findings were generally in lab settings, and it remains unclear how integrating LLM tools into CS1 course instruction affects student learning. This study aims to address that gap.

2.2 Differences in Student Outcomes in CS1

CS1 courses often do not equally support all students. In the United States—where our course takes place—women, Black, Latine/x/a/o, and Native American students are often underrepresented in computing [7]. Demographic factors and pre-college educational access can influence CS1 outcomes, making it crucial to address structures that may hinder students from diverse backgrounds. This section reviews literature highlighting these disparities to ensure that new courses, like our LLM integrated CS1 course, are designed to reduce or not worsen these gaps.

¹<https://techcrunch.com/2025/07/30/github-copilot-crosses-20-million-all-time-users/>

2.2.1 Prior Experience. Although CS1 is intended for beginners, it often favors students with prior programming experience, who consistently outperform their peers [19]. This dynamic can discourage less experienced students [36], and we do not yet know how LLMs might affect this pattern.

2.2.2 English as a Second Language. English proficiency influences CS1 success, as programming instruction and code are primarily in English [17]. Non-native speakers often face language-based challenges that can impact performance [37]. As a community, we need to understand whether and how such challenges present when working with LLM tools.

2.2.3 Socioeconomic Status. Students from higher SES backgrounds tend to perform better in CS1 [51], partly due to related factors like prior experience and parental education [34, 46]. These advantages convey access to resources and academic norms that can widen equity gaps. Curriculum design must consider and address these disparities.

2.2.4 Race and Gender. Racial disparities in CS1 outcomes are linked to differences in pre-college access and systemic biases in instruction [20, 32]. While there are inconsistent findings related to gender-based performance gaps [6, 50], factors like self-efficacy and belonging strongly influence outcomes [23, 28]. We examine how integrating LLMs might influence these dynamics.

2.3 CS1 Learning Objectives in the Age of LLMs

With increasingly capable LLMs, instructors must reconsider which skills remain critical and which new ones should be taught in introductory programming. Prompting is one such emerging skill, requiring students to articulate a program’s purpose in natural language. This natural language reasoning parallels “Explain in Plain English” (EiPE) questions, where students describe code behavior in words; studies suggest that EiPE questions can be used to develop code comprehension and prompting abilities in tandem [12, 49].

While EiPE has long been included in traditional (non-LLM) CS1 contexts, other skills—particularly debugging and testing—have taken on new importance in the LLM era. Despite LLMs’ capabilities, Dakhel et al. [10] found that Copilot can still produce buggy code, emphasizing the importance of students’ mastery of programming fundamentals to evaluate LLM output. Identifying such bugs relies on testing practices, which some instructors are bringing to their CS1 courses [14]. Interviews with instructors [25] further indicate this shift toward new skills in courses that teach programming with LLM tools. This study’s course design addresses these needs.

3 Course Design

This work builds on our earlier report of our GenAI-integrated CS1 course [52]. Here, we include our guiding principles and their justification (Section 3.1), the learning objectives derived from these principles (Section 3.2), and our course implementation choices (Section 3.3).

3.1 Design Principles for CS1-LLM

Here we summarize our design principles (Figure 1). Importantly, they are *complementary* to existing best practices in CS1.

Design Principle 1 – Enable and Encourage Use of LLMs: Throughout the course, instruction and assessments were designed to

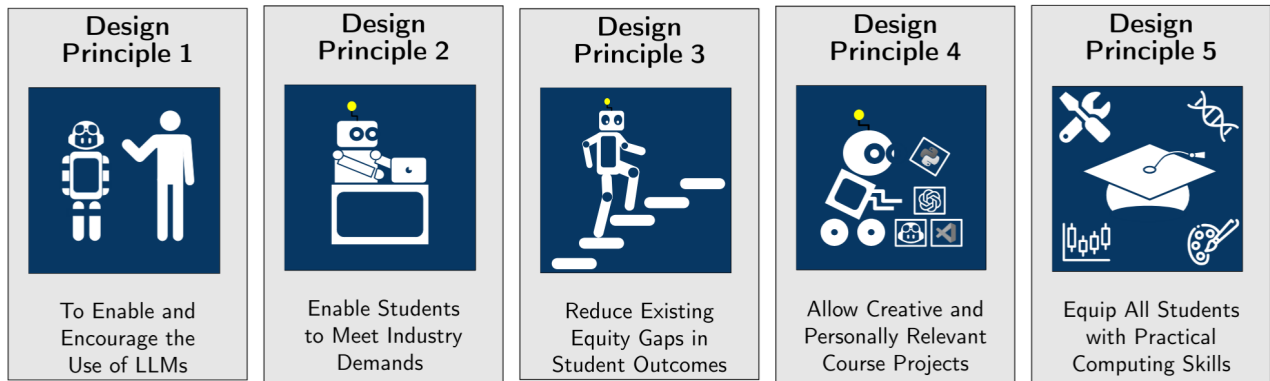


Figure 1: Design Principles for CS1-LLM

include LLM-integrated activities emphasizing skills like open-ended design, prompting, reading, and evaluating LLM-generated code [40]. While students in LLM-integrated environments will be able to generate code without strong programming fundamentals, they will need the aforementioned skills in order to work with this code that they have generated.

Design Principle 2 – Enable Students to Meet Industry Demands: There is often a gap between what traditional computer science coursework prepares students for and the skills required in industry [8]. While debate continues about how much CS1 should emphasize industry preparation, students must still master tools and concepts needed for their majors and careers. Our approach emphasizes code reading, problem decomposition, and testing—skills critical as students work with and assess LLM-generated code that may initially contain errors [10]. These skills are reinforced through both homework and non-LLM exam questions to ensure that students are building these skills independently, and through project-based work that encourages the use of LLMs.

Design Principle 3 – Reduce Equity Gaps in Student Outcomes: Making CS1 courses more equitable across underrepresented groups has been a central topic of research within computer science education. Despite best practices that help address performance gaps between under- and highly-represented groups [9, 54], these gaps persist [42, 43]. Student-to-student interaction practices such as pair programming and peer instruction have also been shown to make CS1 courses more equitable [9, 54]. Our CS1 course—both before and after the incorporation of LLM tools—employs a number of these techniques including peer instruction and contextualized computing. It additionally includes open-ended projects that promote student agency and choice; we suspected that LLMs would be particularly beneficial on such large tasks.

Design Principle 4 – Allow Creative and Personally Relevant Course Projects: Building on *Design Principles 2 and 3*, we designed projects that reflect real-world, project-based work and allow students to explore creative, personally meaningful topics. Many real-world contexts involve larger, less-structured tasks than the Many Small Programs (MSPs) approach of traditional CS1 homework questions [3]. We use the One Large Program (OLP) model [2], enabling deeper integration of skills and authentic applications. LLMs help reduce

the barrier to completing OLPs, allowing students to engage with realistic problem decomposition and pursue personalized goals.

Design Principle 5 – Equip All Students With Practical Computing Skills: As computing spreads across disciplines, accessible introductory courses are essential. Many address this through domain-specific applications such as “Media Computation” [18], where students manipulate images and sounds. Beyond CS1, CS+X programs help students apply computing within other fields [47]. LLMs offer a new way to teach practical computing, helping students tackle advanced, relevant tasks early without traditional programming mastery.

3.2 Learning Objectives

Even before the advent of LLMs led to tectonic shifts in computing education, there was little consensus on the central learning objectives of an introductory computing course [5]. The appropriateness of objectives depends on the student population for which those outcomes are calibrated.

We adapted our course to include specific learning objectives related to the use of LLMs in programming and to de-emphasize traditional learning objectives focused on writing code from scratch. Although less of a focus, our learning objectives still include code writing from scratch for two primary reasons. One is that we wish for our students to be prepared for other contexts in which they may be expected to code without use of GenAI (such as upper level courses that do not use GenAI). Another is that the importance of code writing as a foundational learning process is not yet clear in the emergent GenAI programming landscape. We appreciate the need for further studies to establish the importance of code writing in a student’s journey toward programming proficiency, and therefore continue to have it as a learning objective in our course, though we now emphasize practice in other domains as described below. We organize our objectives using Bloom’s taxonomy [22]:

Level 1: Knowledge Students will be able to define: nondeterminism, Large Language Model (LLM), prompt, prompt engineering, code correctness, problem decomposition, and top-down design.

Level 2: Comprehension Students will illustrate the workflow used when programming with an AI assistant and describe the purpose of common Python programming features, including: 1)

Table 1: Course Schedule

Week	Topic(s)
1	Functions and Working with Copilot
2	Variables, Conditionals, Memory Models
3	Loops, Strings, Testing, VSCode Debugger
4	Loops, Lists, Files, Problem Decomposition
5	Intro to Data Science, Dictionaries
6	Revisit Problem Decomposition and Testing
7	Intro to Images, PIL, Image Filters
8	Copying Images, Intro to Games and Randomness
9	Large Game Example
10	Python Modules and Automating Tedious Tasks

Table 2: Course Assessments and Weights

Assessment Type	Component	Weight
Formative (35%)	Peer Instruction Participation	5%
	Reading Quizzes	5%
	Homework	15%
	Labs	10%
Summative (65%)	Projects	10%
	Mid Quarter Exams	30%
	Final Exam	25%

Variables, 2) Conditionals, 3) Loops, 4) Functions, 5) Lists, 6) Dictionaries, and 7) Modules.

Level 3: Application Students will apply prompt engineering to influence code generated by an AI assistant.

Level 4: Analysis Students will analyze and trace a Python program to determine or explain its behavior, divide a programming problem into subproblems as part of top-down design, and debug a Python program to locate bugs.

Level 5: Synthesis Students will design open- and closed-box tests to determine whether code is correct, identify and fix bugs in Python code, modify Python code to perform a different task, and write complete and correct Python programs using top-down design, prompting, testing, and debugging.

Level 6: Evaluation Students will judge whether a program is correct using evidence from testing and debugging.

3.3 Course Implementation

Our new CS1 course takes place at a large research-focused university in the United States. The course schedule and weighting of assignments appear in Tables 1 and 2. The LLM environment introduced was GitHub Copilot in Microsoft’s Visual Studio Code, selected because of their seamless integration and accessibility via the GitHub Student Developer Pack (free of charge for students). Below we expand on the course components.

Textbook: Our teaching staff used the textbook “Learn AI-Assisted Python Programming with GitHub Copilot and ChatGPT”, assigning weekly readings and graded quizzes. The book combines traditional CS1 topics with LLM-related material like effective Copilot usage and prompt engineering [39].

Lectures: Lectures included mini-lectures, Live Coding [44], and Peer Instruction [38], covering typical CS1 topics along with discussion of Copilot interactions. Instructors highlighted incorrect Copilot responses, prompting students to identify issues, and demonstrated Copilot Chat in VS Code for answering code-related questions.

Formative Assessments: Our course used PrairieLearn for homework and assessments, leveraging features like problem variants and autograding [55]. Homework included multiple choice, short answer, Explain in Plain English [33], Parsons Problems [35], debugging, and code writing. Students were encouraged to work without Copilot for most problems (in preparation for summative assessments) but could use it if needed. Some coding questions were intended for LLM usage, featuring a Copilot-enabled workspace.

Summative Assessments: The course included four 50-minute mid-quarter exams featuring code tracing, explaining [53], Parsons Problems [35], and small coding tasks. Three exams excluded Copilot, but one question on the second exam allowed Copilot access. We originally intended to include Copilot sections on all exams following Exam 1, but technical issues arose. The login process for the Exam 2 Copilot question took longer than expected because of the need for students’ phone access to complete two-factor authentication, as well as some students forgetting their credentials. This led us to forgo Copilot portions on Exams 3 and 4, and only include Copilot again on the final where more time was available.

The final exam consisted of three parts, all of which were proctored:

- (1) **Multiple-choice** (90 min, 70%) — tracing, explaining, testing, debugging, Parsons, and short-answer items.
- (2) **Non-Copilot coding** (45 min, 15%) — four coding questions in PrairieLearn on lab computers.
- (3) **Copilot-assisted coding** (45 min, 15%) — one large problem on personal or lab computers, with sample test cases and partial credit. Variants involved either analyzing a new dataset or creating a restricted spell-check tool.

4 Analysis and Results

Our analysis includes comparing student performance in our course to international CS1 benchmarks (Section 4.1), examining students’ perceptions of programming with LLMs (Section 4.2), and analyzing performance differences across demographics (Section 4.3).


4.1 RQ1) LLM-integrated CS1 vs CS1 Grades

In this section, we compare our course’s student performance to international CS1 course averages reported in Simon et al. [45] to assess our students’ understanding of programming fundamentals while learning with LLMs.

4.1.1 Data Collection and Analysis. On the final exam, we included questions from the Simon et al. [45] study (Table 3). With the exception of Q7, a short answer “Explain in Plain English” question, all questions were multiple-choice or fill in the blank. We graded Q7 out of 4 points with a rubric based on SOLO taxonomy categories [30]. We awarded full points for correct relational responses, 3 points for correct but less abstract responses, 2 points for responses that explained the code line by line, and 1 or 0 points for serious misunderstandings.

Table 3: Questions used from Simon et al. [45]

Q#	Topic	Task	Question Type	Skill Required
Q1	Boolean Expressions	Trace a boolean expression with relational and logical operators.	Multiple Choice	Trace code
Q2	Assignment and Sequence	Trace simple assignment and sequence.	Multiple Choice	Trace code
Q3	Conditional	Trace a nested if statement.	Short Response	Trace code
Q4	Conditional	Trace a nested if/elif to find a variable's value.	Multiple Choice	Trace code
Q5	Iteration and Arrays	Trace a while loop with conditionals and arrays.	Short Response	Trace code
Q6	Iteration and Arrays	Determine the outcome of a for loop over an array.	Multiple Choice	Explain code
Q7	Iteration	Describe the purpose of loop-based code.	Short Response	Explain code

**Figure 2: Simon et al. [45] Benchmarking Data Results**

In total, we had 535 responses to all questions except Q7 ($n = 250$) for which we had to remove the first cohort of test-taker responses due to an error in the question. We compare the performance of students in our course to that of the international averages presented in the Simon et al. [45] study.

4.1.2 Results. Overall, the results of the question level comparisons suggest that our students' learning with LLM tools was largely comparable to past students' learning without these tools. Students in our LLM-integrated course performed similarly to those from historical international averages (Figure 2). The one question type in which students in our LLM-integrated course performed notably better (11%) than the international average was on the short answer "Explain in Plain English" question (Q7). This may suggest that students in our course were better equipped to understand and articulate the purpose of code, perhaps as a result of practicing how to prompt LLMs and evaluate their output.

4.2 RQ2) Students' Perceptions of CS1-LLM

In this section we explore students' perceptions of learning to code with Copilot.

4.2.1 Data Collection and Analysis. To gain an understanding of student perceptions of the course, we distributed two surveys, one midway through the quarter (*mid*) and the other at the end of the quarter prior to the final exam (*end*). The *mid* and *end* surveys asked students a variety of Likert questions along with the following open-ended question:

"How did you feel about working with Copilot as you learned to program this quarter?"

Two researchers qualitatively analyzed responses to this question from both the mid- (400 respondents) and the end-of quarter surveys (315 respondents). To build a code book, they independently

inductively coded the mid-survey responses in batches of 100 and used negotiated agreement to resolve disagreements. After reaching saturation with no new codes emerging, they then independently coded an additional 100 responses to establish inter-rater reliability ($\alpha = 0.86$) using Krippendorff's Alpha [24]. The remaining mid- and end-of-quarter survey responses were then deductively coded using the code book. We present the themes that emerged from this process and the codes that underlie those themes.

4.2.2 Results. We highlight four themes that emerged from the qualitative analysis: Theme 1) *Copilot is helpful for supporting novices to complete programming tasks*, Theme 2) *Copilot can be helpful for learning*, Theme 3) *Some difficulties emerge when working with Copilot*, and Theme 4) *Students can feel overreliant on Copilot*.

Theme 1. Copilot is helpful for supporting novices to complete programming tasks. Students were largely positive regarding programming with Copilot, with 69.2% of mid-quarter survey respondents and 42.2% of end-quarter survey respondents ($N_{mid} = 272$, $N_{end} = 133$) mentioning that Copilot was helpful. Several students indicated that it *saved time* ($N_{mid} = 21$ [5.3%], $N_{end} = 6$ [1.9%]), and *helped them fix errors* ($N_{mid} = 15$ [3.8%], $N_{end} = 5$ [1.6%]):

"I feel that Copilot is very useful in terms of expediting tasks when coding, and also playing the role of an aid when not knowing what to do.— P 79 | MID

"I feel like copilot was a great tool that was able to assist me. Whenever I'm developing code and feel like I ran into a bug, I use copilot to help debug and it saves me so much time on my projects and assignments.— P 531 | MID

One way that students specified that Copilot affords coding support was through *help with questions/hints* about code ($N_{mid} = 32$ [8.1%], $N_{end} = 18$ [5.7%]). Students reported finding it helpful to ask Copilot questions regarding the functionality and syntax of code. Additionally, some students found Copilot helpful for *overcoming obstacles* such as not knowing where to get started, overcoming some form of "coding writer's block", or more generally, when they were simply stuck while coding ($N_{mid} = 37$ [9.4%], $N_{end} = 19$ [6%]):

"Copilot was extremely helpful and when I was stuck it would greatly help when it came to explaining code.— P 415 | END

In particular, several students indicated that they found Copilot particularly *helpful as individuals with no prior programming experience*, as it helped them navigate coding for the first time ($N_{mid} = 12$ [3.1%], $N_{end} = 4$ [1.3%]).

“Programming is a completely new concept to me so I feel as though Copilot was a good helping hand to help get my toes in the water and experience the coding world for a first timer.— P 429 | MID

Students found Copilot helpful for their programming experience in a variety of ways including ideating about code, understanding it, and fixing it. These features may have been especially helpful in assisting novice programmers to get started with coding. This suggests that one of Copilot’s key benefits, from students’ perspectives, is its ability to reduce friction in expressing their ideas in code.

Theme 2. Copilot can be helpful for learning. In reflecting on the connection between their learning process and the use of Copilot, many students ($N_{mid} = 45$ [11.5%], $N_{end} = 26$ [8.3%]) reported that Copilot was *effective for learning and/or understanding code or course material*:

“It helped to understand more complicated problems that I would not have been able to do myself. After reading through the generated code, I get a better understanding of how to solve the problem.— P 366 | END

“I liked Copilot because it was handy when I needed help understanding how to do something or when I was lost I would then look at the code produced and work backwards.— P 415 | MID

Students found Copilot’s outputs to be helpful in understanding how to approach problems, and in understanding how code functions. Given these encouraging findings, it may be worthwhile to explore the particular code-related topics that students found Copilot useful for, the specifics of how they used it, and how these findings can translate to curricular or tool design.

A small subset of students specified that they found unfamiliar code that Copilot generated to be helpful by allowing them to learn from new techniques ($N_{mid} = 19$ [4.8%], $N_{end} = 6$ [1.9%]). In particular, students stated that they were able to learn functions and syntax with which they were previously unfamiliar.

“Working with Copilot has been great for helping me understand how to format my code, adjust errors, and learn about new coding techniques.— P 479 | END

“I feel as though it was an amazing tool that helped me in situations that I felt stuck in. It also helped me to diversify my coding syntax which helped the coding process go by more smoothly.— P 5 | MID

Our findings suggest that some students find these new, unfamiliar code constructs helpful, and instructors may consider ways to intentionally introduce students to these more advanced constructs that LLMs often generate.

Theme 3. Some difficulties emerge when working with Copilot. While several students found it helpful to learn from new code in Copilot’s outputs as described in Theme 2, some did find its unfamiliar outputs to be confusing. A small number of students found it *difficult to interpret output* ($N_{mid} = 6$ [1.5%], $N_{end} = 4$ [1.3%]), which sometimes led to difficulties when the output was incorrect and *needed to be corrected* ($N_{mid} = 7$ [1.8%], $N_{end} = 10$ [3.2%]):

“I feel like Copilot has been helpful, however, because of my lack in programming knowledge, I sometimes struggle with understanding the complex code that Copilot had generated.— P 90 | MID

“It does help giving me a start but I need to modify it a lot sometimes to get to what I want.— P 449 | END

As LLMs may provide code that uses advanced features of the language or code that has bugs, it is not surprising that some students reported being challenged by encountering such code. We recommend educators warn students that this can happen and offer suggestions on how to respond when it does. For example, educators can encourage students to use Copilot to explain the unfamiliar code.

Theme 4. Students can feel overreliant on Copilot. Although most students felt that Copilot was helpful for learning as described in Theme 2, many students expressed concerns about being *overreliant on Copilot* for solving coding problems ($N_{mid} = 94$ [23.9%], $N_{end} = 80$ [25.4%]), stating that they used Copilot more than they felt that they should while coding. Many also mentioned that they felt that it was *hindering their learning* of the course content and fundamentals of programming ($N_{mid} = 59$ [15%], $N_{end} = 44$ [14%]):

“By using Copilot to program, I think that it interfered with my ability to code independently because Copilot did all the work for me. Thus, I struggle to code without Copilot.— P 29 | END

“Copilot definitely helped me get through assignments, but I feel like I relied on it too much and now lack a fundamental understanding of programming basics. If I were asked to code without Copilot, I wouldn’t feel very confident in myself despite doing well in the course.— P 44 | END

Students’ feelings surrounding overreliance were evident in several comments that *Copilot should be introduced later* ($N_{mid} = 15$ [3.8%], $N_{end} = 8$ [2.5%]):

“I feel like there are too many assignments that require Copilot too early on that I become dependent on Copilot. I think the first few weeks should be learning to program without Copilot and later for projects we should use Copilot in a limited capacity to help us program complicated functions. I feel like I just ask Copilot and it gives me the answer but I don’t program a lot without it which is a problem.— P 282 | MID

While some students felt that the course should introduce Copilot later, some addressed their dependence on Copilot by *self-regulating their Copilot use*, often through decreasing the amount of time that they spent using Copilot. Reports of these self-regulation behaviors emerged in our end-of-quarter survey ($N_{end} = 13$ [4.1%]):

“I turned it off after week 2/3, it should be introduced once the first project rolls around, I know people who can’t code because of co-pilot.— P 478 | END

“At first, using Copilot interfered with me learning how to program but as I began to use it in more moderation such as attempting to problem solve and work out coding on my own first and resorting to Copilot only when I am really really stuck on a problem has helped.— P 153 | END

If I had to guess, professional programmers use Copilot or similar tools...




Figure 3: Students' perceptions of how often professionals use LLM tools

This raises questions of what “over-reliance” on Copilot means to students. Familiarity with LLM tools was a core learning objective of the course, and so we wonder why students felt negatively about their reliance on Copilot. It would be strange, for instance, if students stated that they felt over-reliant on a debugger after having been introduced to and encouraged to use a debugging tool. We posit two possible reasons for students feeling as if they were over-reliant on Copilot: 1) lack of access during summative assessments and 2) beliefs about use of Copilot by professional programmers:

Lack of Access During Summative Assessments: One factor that likely affected how reliant students felt they should be on Copilot was the instructor decision to restrict access to Copilot during the majority of their summative assessments. The lack of access on exams may have signaled to students a value in completing certain tasks without the use of Copilot, potentially causing students to see Copilot as a crutch more than a resource. This is reflected from the several students who expressed *summative preparation concerns* ($N_{mid} = 16$ [4.1%], $N_{end} = 6$ [1.9%]):

“Although it was helpful, it was really difficult for me on exams, the whole course is based on using AI tools such as copilot for help, yet **taking it away on exams when we had it in every other type of work seemed a little unfair.**— P 176|END

“Its decent, and helpful- but this class takes it too far. **We need to be able to code without it, OR we can't be quizzed without it. OR make assignments that are specifically easier and designed to be done without Copilot.**— P 484|END

These comments highlight the difference in access between formative and summative assessments. As set forth in the course design, the teaching staff aimed to ensure that students were able to meet core competencies both with and without access to generative AI. However, limiting access to Copilot on summative exams may have impacted student perceptions of the role of LLMs, perhaps causing student overreliance beliefs.

Beliefs about Use of Copilot by Professional Programmers: Another factor that could contribute to feelings of overreliance is that students may not see Copilot as an industry-standard tool, but instead as a tool that enables them to complete assignments with or without what they consider to be a strong understanding of course material. We found that 43.8% of students thought that professional programmers only use Copilot or similar tools rarely or sometimes, while the other 56.2% thought that professionals use these tools routinely or every day (Figure 3). It is possible that with almost half of the class thinking that professional programmers do not routinely use Copilot, frequent use of—or reliance on—such tools was not seen

as acceptable. These students may consider Copilot to be a crutch that professional programmers do not or should not need, rather than a legitimate part of emerging software engineering practices.

4.2.3 Mixed Perceptions about Copilot: Concurrent Themes. In open ended responses, a subset of the class reported a tension between the helpfulness of Copilot as a tool for learning and its potential to hinder learning by always providing an immediate answer when they encountered difficulties ($N_{end} = 67$ [21.3%]):

“I feel a bit conflicted, since it has **helped me write code and understand functions at times**, but I also feel that **I didn't apply myself very well to the course and used Copilot as a crutch when using it to write code and explain code.** Overall, I liked working with Copilot, but my work habits caused me to rely on it too much sometimes.— P 25|END

The sentiments of Copilot being both helpful and hindresome to learning reflect the tension that exists between helping students gain proficiency with LLM usage and providing students the foundational skills needed to code independently.

4.3 RQ3) Demographic Outcomes

In this section, we evaluate how students in different demographic groups in our course compare in their performance outcomes. Some gaps in outcomes between demographic groups have been studied extensively in the CS education community, as described in Section 2.2, and here we explore how outcomes in our LLM-integrated course may differ from historical trends.

4.3.1 Data Collection and Analysis. We focus on exam and project performance outcomes:

- (1) **Average summative exam grade:** The average of all four mid-quarter exams and the final exam.
- (2) **Average project grade:** The average of the three open-ended projects given to students throughout the semester.

See Figure 4 for the distribution of scores.

To analyze how these performance outcomes statistically differ, we performed two Ordinary Least Squares (OLS) regressions: once for average exam score and once for average project score.

$$\begin{aligned} \text{Grade} = & \beta_0 + \beta_1 \text{Non-Male} + \beta_2 \text{First Gen} \\ & + \beta_3 \text{Low Income} + \beta_4 \text{Prior Experience} \\ & + \beta_5 \text{ESL} + \beta_6 \text{Computing Major} + \beta_7 \text{BLNPI} \end{aligned}$$

Here, each of the predictor variables was coded as a binary variable for: **Non-Male** whether the student identified as non-male (female or non-binary), **First Generation** if they were a first generation college student, **Low Income** if they were a low-income student as determined by Pell Grant eligibility, **Prior Experience** if they had prior programming experience, **ESL** if they learned English as a second language, **Computing Major** if they were in a computing-focused major, and **BLNPI** if they identified as Black, Latine/x/a/o, Native American, or Pacific Islander. After filtering out students who did not complete the final exam, project, or demographic survey, we were left with a total of 207 students.




Figure 4: Average exam and project grades

Table 4: OLS regression results for average exam grades and average project grades by demographic group

(a) Average Exam Grades

Variable	β	SE	t	p	95% CI
Intercept	88.8785	2.262	39.297	0.000	[84.418, 93.339]
Non-Male (N=119)	-1.9343	1.644	-1.177	0.241	[-5.176, 1.307]
First Gen (N=105)	-2.3014	1.993	-1.155	0.249	[-6.231, 1.628]
Low Income (N=119)	-1.8873	1.919	-0.984	0.327	[-5.671, 1.896]
Prior Experience (N=100)	3.1900	1.610	1.981	0.049*	[0.014, 6.366]
ESL (N=50)	-0.4636	1.835	-0.253	0.801	[-4.082, 3.154]
Computing Major (N=144)	3.3696	1.780	1.893	0.060	[-0.141, 6.880]
BLNPI (N=72)	-8.4021	1.847	-4.550	0.000***	[-12.044, -4.760]
N = 207 Adj. R ² = 0.216 F = 9.107					

Significance levels: *** p < 0.001, ** p < 0.01, * p < 0.05.

(b) Average Project Grades

Variable	β	SE	t	p	95% CI
Intercept	92.3976	3.328	27.762	0.000	[85.835, 98.961]
Non-Male (N=119)	-0.5270	2.419	-0.218	0.828	[-5.297, 4.243]
First Gen (N=105)	-3.6446	2.932	-1.243	0.215	[-9.427, 2.137]
Low Income (N=119)	1.5766	2.823	0.558	0.577	[-3.991, 7.144]
Prior Experience (N=100)	0.5965	2.370	0.252	0.802	[-4.077, 5.270]
ESL (N=50)	-1.1723	2.700	-0.434	0.665	[-6.496, 4.152]
Computing Major (N=114)	0.1832	2.620	0.070	0.944	[-4.983, 5.349]
BLNPI (N=72)	-2.4990	2.717	-0.920	0.359	[-7.858, 2.860]
N = 207 Adj. R ² = -0.011 F = 0.670					

4.3.2 *Grade Outcomes on Exams.* The results of the exam grades regression in Table 4 show that certain demographic groups were more likely to have significantly higher or lower exam scores including BLNPI students ($\beta_7 = -8.40, p < 0.001$) and students with prior experience ($\beta_4 = 3.19, p = 0.049$).

4.3.3 *Grade Outcomes on Projects.* The results of the project grades regression in Table 4 show that performance is not at all predicted by demographics ($R^2_{adj} = -0.011$)—likely because of a ceiling effect on project grades. Given the plentiful high-level scaffolding and access to low-level programming help from Copilot throughout the process of project completion, it may be the case that there was sufficient infrastructure available for students of all backgrounds to successfully create the projects.

5 Discussion

Our study provides an initial examination of student outcomes and perceptions in a CS1 course that fully incorporates LLMs. In summary, our core findings are:

- (1) Students performed similarly to previous offerings of a CS1:** Despite differing course goals and modified learning objectives, students in our course performed comparably to students in previous offerings of a CS1 course on fundamental programming skills such as code tracing and code explaining. This is encouraging for our community’s general concern of students continuing to learn fundamentals of programming.
- (2) Students reported positive experiences with Copilot with some concerns:** While the vast majority of students

reflected positively with respect to productivity, some expressed concerns on the implications of overreliance.

- (3) Some existing equity gaps persist in exam scores:** Performance disparities by race and prior experience remained on exams (Table 4), reflecting historical findings.
- (4) Projects mitigate equity gaps:** Project grades did not reflect the historical equity gaps found in exams (Table 4). We found *no* significant differences between historically underrepresented groups and highly represented groups on projects.

Next, we discuss implications of our findings for educators, curriculum designers, and the broader research community.

5.1 Implications for Educators

Our results suggest the need for *clearer communication to students* on expected mastery of programming without access to Copilot both within the course *and* in the world of professional programming. Helping students calibrate with respect to these expected outcomes is essential in ensuring that they perceive what they learn in the course as being reflective of what “real programmers” both know and do. This is important in establishing and bolstering students’ self-assessments of themselves as “real programmers” and, as such, carrying with them a sense of belonging [23] and self-efficacy [29]. Future work should consider additional mechanisms by which these expectations can be communicated to students and possible alterations to the curriculum that can facilitate this.

Additionally, the interface between CS1 courses and follow-on courses will likely impact overall student success and which subgroups of students succeed. If subsequent courses ignore that the

learning outcomes have changed in CS1, students may unnecessarily struggle on some topics while being overprepared on others. For example, in our course, we de-emphasize writing code from scratch (because we believe that this skill now has a diminished importance); instructors within a department need to communicate to ensure smooth transitions between courses.

Our finding that projects mitigate historical equity gaps is a promising one, and suggests that educators may consider incorporating open-ended projects like ours that allow LLM usage. While exams remain a source of tension for many groups in introductory computing courses, projects may provide an assessment environment in which students from a variety of backgrounds can thrive. In fact, returning to our learning objectives in Section 3.2, projects, rather than exams, may better target our higher-level objectives of problem decomposition, debugging, and writing complete programs.

5.2 Implications for Course Design & Research

In the study by Lau and Guo [25], the authors found that instructors maintained one of two perspectives toward teaching with LLMs. Some instructors wanted to ban LLM tools, feeling that using LLMs would lead to an erosion of “fundamentals”. Other instructors wanted to embrace AI tools for reasons similar to our own: to prepare students for future jobs and teach more advanced material. Our study may offer useful insights into this ongoing ban vs. embrace debate in favor of embracing LLMs—particularly with respect to the more complex and open-ended assignments students were able to complete under this approach.

One open question for the community that was exposed by our study is whether it is important for students to continue to code from scratch in introductory courses now that LLMs can write that code. In our course, roughly half of the grade was for performance on non-LLM exam questions. This reflects the belief that there is still something to be gained by writing code from scratch, perhaps in terms of building student mental models of state and execution. In fact, through our surveys (in particular, see Theme 4 in our qualitative analysis), we discovered that students, too, want to be able to code without help from LLM tools. How much should students code from scratch? Should they write more code early in the term and progressively off-load it to the LLM? These are important questions for the broader research community.

6 Limitations

This study was conducted in one course at one institution. The depth of comparisons we can make between this course and the international average are limited as this institution may under- or over-perform relative to that average. Only 400 students out of 556 completed the mid-quarter survey, 315 completed the end-of-quarter survey, and 207 provided demographic information. Lastly, as this is a new course, comparisons with existing LLM-integrated courses are necessarily limited.

7 Conclusion

This study examines a novel CS1 course integrating LLM tools into introductory programming, developed in response to the growing adoption of LLM tools by students, educators, and professionals. We assessed student performance outcomes as well as student perceptions toward LLM tools. We find that students performed similarly

in our LLM-integrated course compared with earlier groups of students in traditional CS1 courses on questions that involved code reading on shared learning outcomes. The same groups who struggle in traditional CS1 courses struggled on the exams in this course. By contrast, when students were permitted to use LLM tools on large open-ended projects, we saw no such achievement gap. We note that these projects are larger and more complex than those generally assigned to traditional CS1 students, and that they closely match our higher-level course learning outcomes. Student attitudes toward Copilot are mostly positive and indicate that they found the tool helpful, although some expressed feelings of overreliance on it.

Acknowledgments

This work is supported by the National Science Foundation under Grant No. 2417374 and the Google Award for Inclusion Research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or Google Inc. Generative AI tools—Claude and ChatGPT—were used for minor phrasing and wording changes.

References

- [1] Umar Alkafaween, Ibrahim Albluwi, and Paul Denny. 2025. Automating autograd-ing: Large language models as test suite generators for introductory programming. *Journal of Computer Assisted Learning* 41, 1 (2025), e13100.
- [2] Joe Michael Allen. 2021. *A Many Small Programs (MSP) Approach in a CS1 Course*. University of California, Riverside.
- [3] Joe Michael Allen, Frank Vahid, Kelly Downey, and Alex Daniel Edgcomb. 2018. Weekly programs in a CS1 class: Experiences with auto-graded many-small programs (MSP). In *2018 ASEE Annual Conference & Exposition*.
- [4] Brett A Becker, Paul Denny, James Finnie-Ansley, Andrew Luxton-Reilly, James Prather, and Eddie Antonio Santos. 2023. 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*. 500–506.
- [5] Brett A Becker and Thomas Fitzpatrick. 2019. What do CS1 syllabi reveal about our expectations of introductory programming students?. In *Proceedings of the 50th ACM technical symposium on computer science education*. 1011–1017.
- [6] Sylvia Beyer. 2014. Why are women underrepresented in Computer Science? Gender differences in stereotypes, self-efficacy, values, and interests and predictors of future CS course-taking and grades. *Computer Science Education* 24, 2-3 (2014), 153–192.
- [7] Tracy Camp, W Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben. 2017. Generation CS: The challenges of and responses to the enrollment surge. *ACM Inroads* 8, 4 (2017), 59–65.
- [8] Michelle Craig, Phill Conrad, Dylan Lynch, Natasha Lee, and Laura Anthony. 2018. Listening to early career software developers. *J. Comput. Sci. Coll* 33, 4 (2018), 138–149.
- [9] Catherine H Crouch and Eric Mazur. 2001. Peer Instruction: Ten years of experience and results. *American journal of physics* 69, 9 (2001), 970–977.
- [10] Arghavan Moradi Dakhel, Vahid Majdinasab, Amin Nikanjam, Foutse Khomh, Michel C Desmarais, and Zhen Ming Jack Jiang. 2023. GitHub Copilot AI pair programmer: Asset or liability? *Journal of Systems and Software* 203 (2023), 111734.
- [11] Paul Denny, Juho Leinonen, James Prather, Andrew Luxton-Reilly, Thezyrie Amarouche, Brett A Becker, and Brent N Reeves. 2024. Prompt Problems: A new programming exercise for the generative AI era. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*. 296–302.
- [12] Paul Denny, David H Smith IV, Max Fowler, James Prather, Brett A Becker, and Juho Leinonen. 2024. Explaining code with a purpose: An integrated approach for developing code comprehension and prompting skills. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1*. 283–289.
- [13] Tony Haoran Feng, Andrew Luxton-Reilly, Burkhard C Wünsche, and Paul Denny. 2025. From Automation to Cognition: Redefining the Roles of Educators and Generative AI in Computing Education. In *Proceedings of the 27th Australasian Computing Education Conference*. 164–171.
- [14] Amanda S Fernandez and Kimberly A Cornell. 2024. CS1 with a side of AI: Teaching software verification for secure code in the era of generative AI. In

- Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*. 345–351.
- [15] James Finnie-Ansley, Paul Denny, Brett A Becker, Andrew Luxton-Reilly, and James Prather. 2022. The robots are coming: Exploring the implications of OpenAI Codex on introductory programming. In *Proceedings of the 24th Australasian Computing Education Conference*. 10–19.
 - [16] Simone Grassini. 2023. Shaping the future of education: exploring the potential and consequences of AI and ChatGPT in educational settings. *Education Sciences* 13, 7 (2023), 692.
 - [17] Philip J Guo. 2018. Non-native English speakers learning computer programming: Barriers, desires, and design opportunities. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14.
 - [18] Mark Guzdial. 2003. A media computation course for non-majors. In *Proceedings of the 8th annual conference on Innovation and technology in computer science education*. 104–108.
 - [19] Dianne Hagan and Selby Markham. 2000. Does it help to have some programming experience before beginning a computing degree program?. In *Proceedings of the 5th annual SIGCSE/SIGCUE ITiCSE Conference on Innovation and technology in computer science education*. 25–28.
 - [20] Sandra Katz, John Aronis, David Allbritton, Christine Wilson, and Mary Lou Soffa. 2003. Gender and race in predicting achievement in computer science. *IEEE Technology and Society Magazine* 22, 3 (2003), 20–27.
 - [21] Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Zachary Henley, Paul Denny, Michelle Craig, and Tovi Grossman. 2024. CodeAid: Evaluating a classroom deployment of an LLM-based programming assistant that balances student and educator needs. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–20.
 - [22] David R Krathwohl. 2002. A revision of Bloom’s taxonomy: An overview. *Theory into practice* 41, 4 (2002), 212–218.
 - [23] Sophia Krause-Levy, William G Griswold, Leo Porter, and Christine Alvarado. 2021. The relationship between sense of belonging and student outcomes in CS1 and beyond. In *Proceedings of the 17th ACM Conference on International Computing Education Research*. 29–41.
 - [24] Klaus Krippendorff. 2011. Computing Krippendorff’s Alpha-Reliability. <https://api.semanticscholar.org/CorpusID:59901023>
 - [25] Sam Lau and Philip Guo. 2023. From “Ban it till we understand it” to “Resistance is futile”: How university programming instructors plan to adapt as more students use AI code generation and explanation tools such as ChatGPT and GitHub Copilot. In *Proceedings of the 2023 ACM Conference on International Computing Education Research-Volume 1*. 106–121.
 - [26] Mark Liffiton, Brad E Sheese, Jaromir Savelka, and Paul Denny. 2023. CodeHelp: Using large language models with guardrails for scalable support in programming classes. In *Proceedings of the 23rd Koli Calling International Conference on Computing Education Research*. 1–11.
 - [27] Jintao Ling and Muhammad Afzaal. 2024. Automatic question-answer pairs generation using pre-trained large language models in higher education. *Computers and Education: Artificial Intelligence* 6 (2024), 100252.
 - [28] Alex Lishinski and Joshua Rosenberg. 2021. All the pieces matter: The relationship of momentary self-efficacy and affective experiences with CS1 achievement and interest in computing. In *Proceedings of the 17th ACM Conference on International Computing Education Research*. 252–265.
 - [29] Alex Lishinski, Aman Yadav, Jon Good, and Richard Enbody. 2016. Learning to program: Gender differences and interactive effects of students’ motivation, goals, and self-efficacy on performance. In *Proceedings of the 2016 ACM conference on international computing education research*. 211–220.
 - [30] Raymond Lister, Beth Simon, Errol Thompson, Jacqueline L Whalley, and Christine Prasad. 2006. Not seeing the forest for the trees: novice programmers and the SOLO taxonomy. *ACM SIGCSE Bulletin* 38, 3 (2006), 118–122.
 - [31] Chung Kwan Lo. 2023. What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences* 13, 4 (2023), 410.
 - [32] Jane Margolis, Rachel Estrella, Joanna Goode, Jennifer Jellison Holme, and Kimberly Nao. 2008. *Stuck in the Shallow End: Education, Race, and Computing*. The MIT Press.
 - [33] Laurie Murphy, Renée McCauley, and Sue Fitzgerald. 2012. “Explain in plain English” questions: implications for teaching. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education*. 385–390.
 - [34] Barbara Martinez Neda, Flor Morales, Kitana Carbajal Juarez, Jennifer Wong-Ma, and Sergio Gago-Masague. 2024. Investigating the Role of Socioeconomic Factors on CS1 Performance. In *2024 IEEE Global Engineering Education Conference*. 1–8.
 - [35] Dale Parsons and Patricia Haden. 2006. Parson’s programming puzzles: a fun and effective learning tool for first programming courses. In *Proceedings of the 8th Australasian Conference on Computing Education-Volume 52*. 157–163.
 - [36] Andrew Petersen, Michelle Craig, Jennifer Campbell, and Anya Taffioch. 2016. Revisiting why students drop CS1. In *Proceedings of the 16th Koli Calling International Conference on Computing Education Research*. 71–80.
 - [37] Nelishia Pillay and Vikash R. Jugoo. 2005. An investigation into student characteristics affecting novice programming performance. *SIGCSE Bull.* 37, 4 (dec 2005), 107–110.
 - [38] Leo Porter, Dennis Bouvier, Quintin Cutts, Scott Grissom, Cynthia Lee, Robert McCartney, Daniel Zingaro, and Beth Simon. 2016. A multi-institutional study of Peer Instruction in introductory computing. In *Proceedings of the 47th ACM Technical Symposium on Computer Science Education*. 358–363.
 - [39] Leo Porter and Daniel Zingaro. 2024. *Learn AI-assisted Python Programming: With GitHub Copilot and ChatGPT*. Manning Publications.
 - [40] James Prather, Paul Denny, Juho Leinonen, Brett A Becker, Ibrahim Albluwi, Michelle Craig, Hieke Keuning, Natalie Kiesler, Tobias Kohn, Andrew Luxton-Reilly, et al. 2023. The robots are here: Navigating the generative AI revolution in computing education. In *Proceedings of the 2023 Working Group Reports on Innovation and Technology in Computer Science Education*. 108–159.
 - [41] James Prather, Juho Leinonen, Natalie Kiesler, Jamie Gorson Benario, Sam Lau, Stephen MacNeil, Narges Norouzi, Simone Opel, Vee Pettit, Leo Porter, Brent N. Reeves, Jaromir Savelka, David H. Smith IV, Sven Strickroth, and Daniel Zingaro. 2025. Beyond the Hype: A Comprehensive Review of Current Trends in Generative AI Research, Teaching Practices, and Tools. In *2024 Working Group Reports on Innovation and Technology in Computer Science Education*. 300–338. <https://doi.org/10.1145/3689187.3709614>
 - [42] Adrian Salguero, William G Griswold, Christine Alvarado, and Leo Porter. 2021. Understanding sources of student struggle in early computer science courses. In *Proceedings of the 17th ACM Conference on International Computing Education Research*. 319–333.
 - [43] Adrian Salguero, Julian McAuley, Beth Simon, and Leo Porter. 2020. A longitudinal evaluation of a best practices CS1. In *Proceedings of the 2020 ACM conference on international computing education research*. 182–193.
 - [44] Ana Selvaraj, Eda Zhang, Leo Porter, and Adalbert Gerald Soosai Raj. 2021. Live coding: A review of the literature. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1*. 164–170.
 - [45] Simon, Judy Sheard, Daryl D’Souza, Peter Klemperer, Leo Porter, Juha Sorva, Martijn Stegeman, and Daniel Zingaro. 2016. Benchmarking introductory programming exams: some preliminary results. In *Proceedings of the 2016 ACM Conference on International Computing Education Research*. 103–111.
 - [46] Selcuk R Sirin. 2005. Socioeconomic status and academic achievement: A meta-analytic review of research. *Review of educational research* 75, 3 (2005), 417–453.
 - [47] Robert H Sloan, Valerie Barr, Heather Bort, Mark Guzdial, Ran Libeskind-Hadas, and Richard Warner. 2020. CS+ X meets CS 1: Strongly themed intro courses. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*. 960–961.
 - [48] David H. Smith IV, Paul Denny, and Max Fowler. 2024. Prompting for Comprehension: Exploring the Intersection of Explain in Plain English Questions and Prompt Writing. In *Proceedings of the Eleventh ACM Conference on Learning@Scale*. 39–50.
 - [49] David H. Smith IV and Craig Zilles. 2024. Evaluating Large Language Model Code Generation as an Autograding Mechanism for “Explain in Plain English” Questions. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education Volume 2*. 1824–1825.
 - [50] Harriet G Taylor and Luegna C Mounfield. 1994. Exploration of the relationship between prior computing experience and gender on success in college computer science. *Journal of educational computing research* 11, 4 (1994), 291–306.
 - [51] Jennifer Alexandra Thompson, Margaret Ellis, and Sara Hooshangi. 2022. High school socioeconomic neighborhood status and CS1 performance. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education Volume 2*. 1348–1348.
 - [52] Annapurna Vadaparty, Daniel Zingaro, David H. Smith IV, Mounika Padala, Christine Alvarado, Jamie Gorson Benario, and Leo Porter. 2024. CS1-LLM: Integrating LLMs into CS1 Instruction. In *Proceedings of the 2024 conference on Innovation and Technology in Computer Science Education V. 1*. 297–303.
 - [53] Anne Venables, Grace Tan, and Raymond Lister. 2009. A closer look at tracing, explaining and code writing skills in the novice programmer. In *Proceedings of the fifth international workshop on Computing education research*. 117–128.
 - [54] Linda L Werner, Brian Hanks, and Charlie McDowell. 2004. Pair-programming helps female computer science students. *Journal on Educational Resources in Computing (JERIC)* 4, 1 (2004), 4–es.
 - [55] Matthew West, Geoffrey L Herman, and Craig Zilles. 2015. PrairieLearn: Mastery-based online problem solving with adaptive scoring and recommendations driven by machine learning. In *2015 ASEE Annual Conference & Exposition*. 26–1238.
 - [56] Yuankai Xue, Hanlin Chen, Gina R. Bai, Robert Tairas, and Yu Huang. 2024. Does ChatGPT Help With Introductory Programming? An Experiment of Students Using ChatGPT in CS1. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering Education and Training*. 331–341. <https://doi.org/10.1145/3639474.3640076>