Evaluating a Course-Specific AI Companion in Undergraduate Computing Courses: A Mixed Methods Pilot Study

Prof. Oguzhan Topsakal, University of South Florida

Oguzhan Topsakal, Ph.D., is Associate Professor of Instruction and Director of Accreditation at the Bellini College of AI, Cybersecurity & Computing, University of South Florida. His work focuses on AI-augmented learning for education, deep learning for medical imaging, and digital surgery planning. He teaches AI, algorithms, automata, and database design. Previously, he was on the Computer Science faculty at Florida Polytechnic University. He holds a Ph.D. and M.S. from the University of Florida and a B.Eng. from Istanbul Technical University.

Jing Wang, University of South Florida
Dr. Schinnel Kylan Small, University of South Florida

Schinnel Small is an Instructor I and IT Director at the University of South Florida. She pursued a Bachelor-to-Doctorate path of study at Morgan State University and obtained her Bachelor of Science and Doctorate degrees in Electrical and Computer Engine

Taseef Rahman, University of South Florida
James Anderson, University of South Florida
William Joseph Gauvin, University of South Florida

I am currently working at Broadcom Inc/Symantec, as a technical director working in the Cyber Security Services/Global Government Engineering (CSS/GGE) group where I work as an architect, designing advanced malware detection and analysis tools using machine learning. I have 43 years of experience in networks and security; having personally been involved in the design and development of 3 firewalls, the AltaVista Firewall, Raptor Firewall and Symantec Gateway Security Firewall/UTM. In addition, I worked on a number of other security related technologies, a version of Norton Antivirus for the Small Mobile Environment / Portable Electronic Device, an advanced version of Data Loss Protection filtering and detection, a secure file reputation system, and Symantec Insight for Malware Researchers (a malware analysis tool). I am highly involved in Cyber Threat Information Sharing (CTIS), for which I was a member of the group that defined the protocols and language (STIX/TAXII) used by the Department of Homeland Security (DHS) to allow government and industries to share cyber threat intelligence. I was also a member of a group that is activity participating in the response to EO 14028, to provide recommendations on how to secure the software supply chain. I have been issued 28 patents on security technologies and received my PhD from the University of Massachusetts; my concentration was in using Machine Learning - Inference of Security Information using Social Networks.

I have been teaching graduate and undergraduate classes since 1988, and teaching is a passion of mine. I am currently a Professor of Practice in Cybersecurity at the University of South Florida, teaching ”Foundations of Cyber Security”, ”Network Security and Firewalls”, ”Network Programming” (at both the undergrad and graduate level) and ”Reverse Engineering and Malware Analysis”. My goal in these courses is to bring my real-world knowledge to the student, using state of the art technologies and hands-on assignments.

Dr. Ellyn Couillard, University of South Florida

Ellyn Couillard is the Director of Student Success for the University of South Florida (USF) Bellini College of Artificial Intelligence, Cybersecurity and Computing. Her research is focused on computing and STEM education. She teaches the first-year seminar, supports academic support services, career development, and student issues within the USF Bellini College. She previously served in administration, leadership, and/or faculty roles at Florida State University, USF Taneja College of Pharmacy, Paine College, and Mercer University. She holds a PhD in Curriculum and Instruction from USF, MS in College Counseling from St. Cloud State University, and BA in Sociology from The University of Georgia.

Robert F. Dedrick, University of South Florida

Robert F. Dedrick is Professor and Coordinator of the Educational Measurement and Research Program in the College of Education at the University of South Florida. He specializes in the development and validation of psychological and educational measures, and teaches courses in research design, measurement, and statistics. His MA and PHD degrees are from the University of Michigan.

Dr. Ken Christensen P.E., University of South Florida

Ken Christensen is Professor and Associate Dean of Academic Affairs in the Bellini College of AI, Cybersecurity and Computing at the University of South Florida. His research interests are at the intersection of networks and energy in the area of energy-efficient computing and communications systems, and application of communications to improve the storage, distribution, and use of energy in small grids. He has over 100 journal and conference publications and 13 U.S. patents. Ken received the PhD in Electrical and Computer Engineering from North Carolina State University in 1991, MS from North Carolina State University in 1983, and BS from the University of Florida in 1981. Ken is an ABET EAC program evaluator. Ken is a licensed Professional Engineer in the state of Florida, a Senior Member of IEEE, and a member of ACM and ASEE

Evaluating a Course-Specific AI Companion in Undergraduate Computing Courses: A Mixed Methods Pilot Study

Oguzhan Topsakal[1], Jing Wang[1], Schinnel K. Small[1], Taseef Rahman[1], James Anderson[1], William Gauvin Jr[1], Ellyn Couillard[1], Robert F. Dedrick[2], Ken Christensen[1]

[1]Bellini College of AI, Cybersecurity and Computing, University of South Florida

[2] College of Education, University of South Florida

Abstract

Purpose-built AI tutors embedded in learning management systems (LMS) may enhance learning and engagement, yet evidence from authentic deployments remains limited. We developed an AI Course Companion tool that grounds answers in instructor-uploaded lecture materials and generates practice questions patterned on prior exams and homework to support targeted studying. Pedagogically, the design implements retrieval practice (exam-style generation), worked-example scaffolding (stepwise explanations), and tight content alignment to reduce extraneous load, consistent with constructivist, cognitive-load, and formative-assessment perspectives.

We evaluated usability and instructional impact of a course-specific, LMS-embedded companion, featuring lecture grounding and prior-assessment-based practice, offered as an optional, ungraded tool in undergraduate computing courses at a large public university. A mixed-methods, one-semester study spanned six courses. Data sources include an end-of-term student survey, an instructor survey, 20-minute interviews on effectiveness and impact, and anonymized usage analytics (feature-level events over time). Primary outcomes were student-reported effectiveness and usability; secondary outcomes were instructor-reported instructional impact (e.g., fewer repetitive questions, alignment); exploratory outcomes were behavioral usage metrics.

Findings indicated broad endorsement alongside clear improvement areas. In recommendation items, 81% of students and 100% of instructors endorsed the tool for future offerings. Students credited course-aligned practice, stepwise explanations, and targeted summaries with improving study efficiency, especially for exams; the most valued attribute was tight alignment to course artifacts (slides, homework, exam style). Pedagogically, it aligned with course objectives and supported asynchronous/online learning by enabling personalized study; with clear guidelines, it augmented rather than replaced student thinking. Requested enhancements included deeper lecture integration, auto-generated practice exams by topic/slide, improved cross-document linking (slides, labs, readings), and a student onboarding assistant; instructors recommended early, planned pre-semester adoption. Interview synthesis underscored that the tool's chief strength was tight alignment with course materials, yielding focused, on-topic responses. It also indicated benefits for proactive students, beginners, repeaters, and regular LMS users. Usage analytics showed increased engagement with the practice-questions feature in the days preceding exams.

Keywords: Computing education, AI in education; learning analytics; intelligent tutoring systems; mixed methods; higher education, AI tutor, AI TA, Generative AI, Education technology, Course-aligned AI Chatbot, Lecture-grounded AI, Course-aligned intelligent system, Conversational AI Agent in Education.

1. Introduction

Large language models (LLMs) have rapidly increased the feasibility of AI tutors, but most empirical evidence still centers on general-purpose chatbots or standalone platforms rather than course-specific companions that are tightly aligned to local instructional artifacts (e.g., lecture slides, assignments, exams). This distinction matters pedagogically: without grounding to course materials, AI systems risk misalignment, inconsistency in explanations, and limited transfer to authentic study tasks. Embedding an AI tutor directly within a learning management system (LMS) further facilitates access, situates help in students' routine workflows, and enables instrumentation of real-world use; which are conditions that are underrepresented in prior evaluations.

We developed an LMS-embedded AI Course Companion designed to operate as a course-specific assistant. The system (a) grounds answers in instructor-uploaded lecture materials, (b) generates practice questions patterned on prior exams and homework, and (c) provides stepwise, worked-example style explanations and concise summaries. These features intentionally implement retrieval practice (exam-style generation), worked-example scaffolding (stepwise solutions), and tight content alignment to reduce extraneous cognitive load. These design choices are consistent with constructivist, cognitive-load, and formative-assessment perspectives. Practically, the tool functions as a "24/7 Teaching Assistant" that is optional, ungraded, anonymous, and available on demand within the LMS.

This pilot study examines the feasibility and perceived value when such a companion is deployed in authentic settings. During a one-semester deployment across six undergraduate computing courses at a large public university (Summer 2025, Bellini College of AI, Cybersecurity and Computing, University of South Florida), we coupled anonymized usage analytics with end-of-term student and instructor feedback. Data sources included an end-of-term student survey (n=139 respondents; analytic subset n=62 students who reported using the tool three or more times) with 11 Likert-type items plus open-ended prompts, an instructor survey (n=5) with Likert items and open-ended prompts, semi-structured instructor interviews (n=3), and feature-level event logs over time. Primary outcomes focused on student-reported effectiveness and usability; secondary outcomes captured instructor-perceived instructional impact (e.g., alignment, reduction in repetitive questions); exploratory outcomes examined behavioral usage patterns.

Guided by this design and context, we addressed four research questions:

- **RQ1 (Student experience):** How do students perceive usability and learning support (e.g., study efficiency, clarity, relevance) of a lecture-grounded, LMS-embedded AI companion?
- **RQ2 (Instructional impact):** How do instructors perceive alignment with course objectives and day-to-day value (e.g., fewer repetitive questions)?
- **RQ3 (Use patterns):** When do students use key features (e.g., practice-question generation), and how do these patterns vary over the term (e.g., before exams)?
- **RQ4 (Design directions):** What strengths and limitations emerge, and what improvements would increase instructional fit and student benefit?

This work makes three contributions. First, it provides empirical evidence from a multi-course, real-world deployment of a course-specific, LMS-embedded, lecture-grounded AI tutor, an under-studied setting relative to general chatbots. Second, it articulates a pedagogical design that operationalizes retrieval practice, worked-example scaffolding, and content alignment inside the LMS, and evaluates that pedagogical design with mixed methods. Third, it elicits concrete, instructor- and student-driven design directions for scalable adoption (e.g., earlier and finer-grained grounding with lecture slide-level citations; structured practice with richer feedback, including topic/slide-level practice exams; improved mobile readability and response speed; cross-document linking; and explicit onboarding).

Findings indicate broad endorsement alongside clear improvement areas: 81% of students and 100% of instructors recommended continued use; students credited course-aligned practice, stepwise explanations, and targeted summaries with improved study efficiency (especially near exams), and instructors reported reduced repetitive questions and strong pedagogical fit when expectations were made explicit. Usage analytics showed increased engagement with practice-question generation in the days preceding exams. The remainder of the paper details the prior work and rationale, the system and context, methods, results, and implications for designing and integrating course-specific AI companions in higher education.

2. Related Work and Rationale

The use of Generative AI has gained significant attention in higher education due to its potential to provide personalized instructional support and increased accessibility. However, the use of Generative AI also raises concerns of plagiarism and trust. It is important for educators to understand the uses of Generative AI in classrooms and design more trustworthy interfaces for education [1].

AI tutors built on large language models (LLMs), including GPT, Gemini, and Llama, are increasingly used to provide scalable, personalized, on-demand academic support across disciplines. These systems can generate natural-language explanations, engage in multi-step reasoning, and converse with students and instructors in ways that approximate human tutoring [2]. An especially promising configuration couples LLMs with course-specific content retrieval from designated knowledge bases (e.g., lecture slides, readings, assignments). By grounding responses in local instructional materials, such tutors can improve alignment with instructor intent and enhance the accuracy and contextual relevance of answers [3].

The AI Course Companion in this study extends these advances with a course-specific tutor designed for targeted university courses. Unlike generic chatbots, it leverages large-context LLM capabilities while explicitly referencing course artifacts to provide explanations that are aligned with stated learning outcomes and course conventions.

Growing evidence suggests that AI tutors can positively influence learning processes and perceived outcomes. Reported benefits include immediate feedback, tailored explanations, and support for varied learner needs [4]. At scale, implementations such as Harvard's CS50 have demonstrated how AI assistance can expand high-touch support in computing courses [5]. By grounding AI responses in official course materials and adopting a "guard railed" approach, the AI-based system in CS50 at Harvard reduces the chances of "hallucination" and improves accuracy in curricular and administrative answers. Other studies in university contexts (e.g., robotics and related computing courses) have examined how advanced LLM techniques can augment lecture-oriented tutoring and improve course engagement [6] and student performance [7].

Instructors also benefit when AI support systems reduce repetitive administrative load; fielding recurring questions or streamlining certain aspects of formative feedback; allowing faculty to dedicate more time to higher-impact instructional work [8, 9]. Another instructor-focused research reveals that instructors recognize the potential of AI tools to assist student learning and new opportunities to improve education with these tools and providing support for both students and teachers [10].

Recent case studies highlight the value of course-specific retrieval of instructional content in improving response relevance and reliability. For example, an AI tutor ("BiWi AI Tutor") that referenced seminar materials reported question-answering accuracy above 84% in its evaluation [11]. Similarly, deployments using Llama-based assistants equipped to reference curricular resources have supported learners in both high-school and college settings, offering real-time guidance that more closely mirrors instructor-like support [12].

Collectively, prior work motivates empirical evaluation of a course-specific AI tutor deployed in authentic classes, with attention to usability, perceived learning value, and instructional integration.

3. Methodology

3.1. Design

We used a mixed-methods design in a one-semester, real-world deployment. A course-specific, LMS-embedded AI Course Companion was introduced at the start of Summer 2025 and remained available as an optional, ungraded resource for the entire term. There was no randomization or experimental manipulation; participation was voluntary and risk was minimal. Quantitative and qualitative strands were analyzed separately and then integrated via convergent synthesis (joint displays and triangulation) to identify corroboration, complementarity, and divergence.

3.2. Setting, Participants, Data Sources

The study took place at the Bellini College of AI, Cybersecurity, and Computing, University of South Florida (USF) in Summer 2025, spanning six undergraduate computing courses taught by five instructors (five courses in-person, face-to-face; one course online, asynchronous):

- COP 2510 – Programming Concepts (F2F)
- COP 4503 – Advanced Database Systems for Information Technology (F2F)
- CDA 3103 – Computer Organization (F2F)
- CAP 4136 – Malware Analysis and Reverse Engineering (Online, asynchronous)
- COT 4210 – Automata Theory and Formal Languages (F2F)
- CDA 3201 – Computer Logic Design (F2F)

Instructors were recruited by email from those scheduled to teach undergraduate courses in Summer 2025. Students were recruited via in-class announcements and Canvas (LMS) announcements in participating courses. Total enrollment was 311 undergraduates across participating courses. The inclusion criteria were being enrolled in a participating course and being over 18 years old. The AI Course Companion was introduced early in the term and accessible within Canvas throughout the semester.

We collected the following data sources and outcomes:

a) Anonymized, feature-level usage analytics over time;
b) An end-of-term student survey (n = 139 respondents; 44.7% of 311) with 11 Likert-type items, a frequency-of-use item, and three open-ended prompts;
c) An instructor survey (n = 5) with six Likert items and six open-ended prompts; and
d) 20-minute semi-structured interviews with three instructors.

Primary outcomes were student-reported usability and perceived learning support. Secondary outcomes were instructor-reported instructional impact (e.g., alignment, reduction in repetitive questions). Exploratory outcomes were behavioral usage patterns (e.g., practice-question engagement near exams). Participation was voluntary with de-identified responses. The activity was conducted as minimal-risk educational research consistent with institutional oversight. The study was approved by the University of South Florida IRB (Protocol STUDY008831).

3.3. Intervention: AI Course Companion

The AI Course Companion (AICC) is a web-based, course-specific AI tutor built on LLMs, designed to answer student questions in natural language by retrieving and referencing course materials (e.g., slides, assignments, readings) for context-aware responses. AICC offers explanations, examples, and feedback aligned with course learning objectives.

AICC is deployed via a web interface and embedded in Canvas as a course page accessible through a course menu item. Use was optional (not tied to grades), and instructors promoted awareness via in-class and LMS announcements. AICC supported anonymous interactions; no identifying data were captured during tool use. Responses were generated using the Google Gemini 2.5 Flash model accessed via API. Guardrails were implemented via prompt-based policies that instructed the model to (a) stay within the scope of course content and (b) decline non-course-related requests and requests for direct solutions to graded assessments.

3.4. Measures

3.4.1. System-Logged Usage

System-logged data captured anonymous usage metrics such as session counts and feature utilization events within the AI Course Companion. To protect participant privacy, all logs were fully anonymized and contained no identifying information that could link usage data to individual students. No personal identifiers were stored during tool use.

3.4.2. Student Survey

Students completed a post-course survey designed to capture their perceptions and experiences with the AI Course Companion. The survey began with a frequency question: “Approximately how often did you use the AI Course Companion this term?” Responses indicating no use or minimal use (only once or twice) were excluded from analysis, leaving an analytic subset of 62 respondents who reported using the tool three or more times during the semester. The survey consisted of both Likert-scale and open-ended questions. Likert items measured ease of use, clarity and relevance of responses, perceived learning support, and overall satisfaction. Sample items included statements such as: “The AI Course Companion was easy to use,” “The tool’s responses were clear and understandable,” and “I would recommend this AI tool for future versions of the course.” Open-ended questions invited participants to elaborate on their experiences, focusing on how the tool supported their learning, any limitations or frustrations they encountered, and suggestions for improving future versions of the AI Course Companion.

3.4.3. Instructor Survey

Instructors who integrated the AI Course Companion into their courses completed a parallel survey that included both Likert-scale and open-ended questions. The Likert items assessed integration with course materials, accuracy and alignment of AI responses, observed student benefits, workload reduction, and willingness to use the tool in future semesters. Open-ended prompts encouraged instructors to describe student feedback, perceived impacts on learning and engagement, challenges or concerns encountered, alignment with pedagogical goals, desired improvements, and overall recommendations for other instructors considering the tool.

3.4.4. Instructor Interview

To gain deeper qualitative insights, instructors participated in semi-structured interviews organized into four thematic sections. The first section, Course Context and Use, explored the course setting, methods of introducing the AI Course Companion, and perceived levels of student engagement. The second section, Perceived Effectiveness, examined instructors’ impressions of the tool’s usefulness, its influence on student engagement and performance, and differences in benefits across student types. The third section, Instructional Impact, focused on how the AI tool affected teaching experiences, including potential time savings, accuracy and alignment concerns, and its compatibility with instructors’ pedagogical goals such as promoting critical thinking and independent learning. The final section, Improvement and Future Use, solicited feedback on desired features, intentions for future adoption, and contexts in which instructors would recommend the tool to colleagues.

4. Results

The pilot study was conducted during the Summer 2025 semester across six courses taught by five instructors at the Bellini College of AI, Cybersecurity, and Computing at the University of South Florida. A total of 311 students were enrolled in these courses, and 139 students completed the post-course survey, representing a 44.7% response rate. Of these respondents, 62 students (44.6%) reported using the AI Course Companion three or more times during the semester. Instructor participation included all five instructors who completed the instructor survey, and three of them also took part in follow-up semi-structured interviews lasting approximately 20 minutes each, conducted via Microsoft Teams and later transcribed for analysis.

Table 1 summarizes the characteristics of the 139 student participants. The majority of respondents were juniors (38.8%) and seniors (38.8%), followed by sophomores (18.7%) and freshmen (3.6%). Most participants were majoring in Computer Science (64.0%), with additional representation from Computer Engineering (15.8%), Cybersecurity (12.2%), Information Technology (7.2%), and Electrical Engineering (0.7%).

Table 1. Student participant characteristics (n = 139)

Characteristic	**n**	**%**
Academic Level		
Freshman	5	3.6
Sophomore	26	18.7
Junior	54	38.8
Senior	54	38.8
Graduate student	0	0.0
Major		
Computer Science	89	64.0
Computer Engineering	22	15.8
Cybersecurity	17	12.2
Information Technology	10	7.2
Electrical Engineering	1	0.7

4.1. System-Logged Usage

System-logged analytics were collected as fully anonymized, feature-level interaction events (with no identifying information) and summarized as daily counts by course. Across the Summer 2025 deployment (May 12–July 25, 2025), the system recorded 2,444 total interactions across the six participating courses. Figure 1 shows daily AI Course Companion interactions by course. Engagement was highly concentrated in COT4210 (Automata Theory), which accounted for 1,589 interactions (65.0%) and showed activity on 61 distinct days, indicating sustained use across the term; the remaining activity was distributed across COP2510 (336; 13.7%), the databases course (logged as COP4703; 233; 9.5%), CDA3103 (151; 6.2%), CDA3201 (102; 4.2%), and CAP4136 (33; 1.4%).

Temporally, interaction volume increased after the start of the semester (May: 187 total interactions) and peaked mid-semester (June: 1,337), remaining elevated through July (920), with several sharp spikes rather than steady daily use (e.g., June 4: 190 interactions, June 11: 168 interactions, July 21: 217 interactions). This “spiky” pattern is consistent with the tool being used as a targeted study aid, with heavier engagement in periods that align with assessments, an interpretation aligned with the study’s focus on behavioral use patterns and the expectation of increased practice-oriented activity near exams.

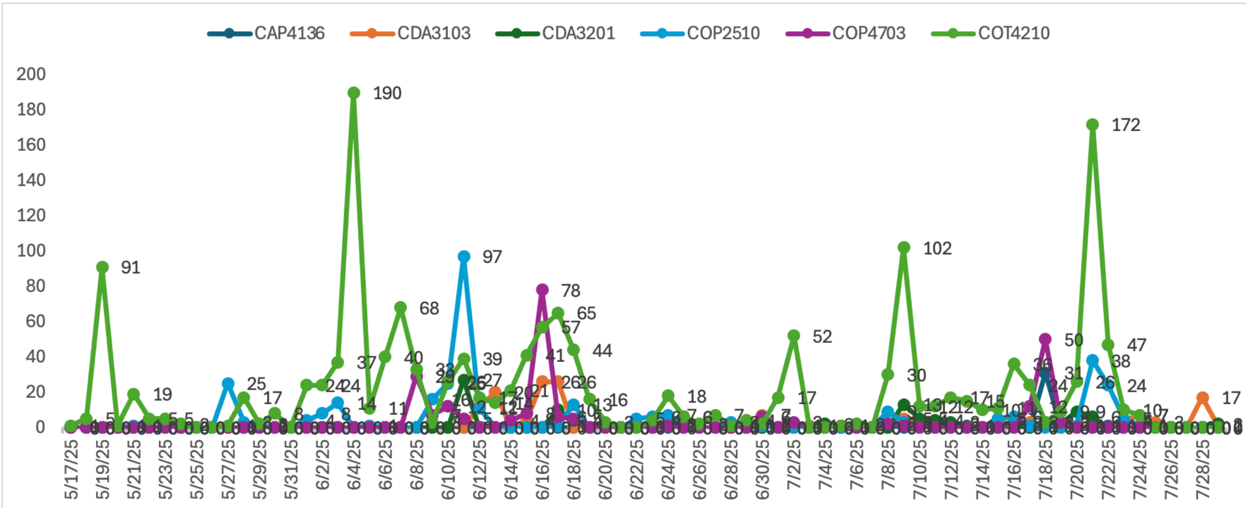

Figure 1: Daily AI Course Companion interactions by course during the Summer 2025 semester

4.2. Student Survey Outcomes

Descriptive analyses of the student survey focused on the 62 respondents who reported using the AI Course Companion at least occasionally ("occasionally," "frequently," or "very frequently") during the term. This analytic subset therefore reflects students with sustained exposure to the system rather than one-time or non-users. Across this group, usage patterns depicted in the frequency chart suggest that the AICC functioned primarily as a recurring study aid that students returned to throughout the semester rather than as a novelty tool used only once near an exam.

Likert-scale items examined perceived usability, clarity and relevance of responses, learning support, and overall satisfaction. The corresponding charts show that responses generally skewed toward the positive end of the scale for students who used the tool at least occasionally. Items targeting ease of use and clarity indicated that most respondents reported the AI Companion straightforward to access and interact with, and that it provided understandable, on-topic answers grounded in course materials. Neutral responses were present but less common than agreement responses, and only a small fraction of students selected clear disagreement categories on these usability items, pointing to a broadly favorable experience with the interface and basic functionality.

Items probing perceived learning support and study impact showed a similar pattern. Across questions about whether the AI Course Companion helped with exam preparation, supported understanding of difficult concepts, and made studying more efficient, the charts indicate that positive ratings outweighed negative ones, with many students endorsing the tool as a helpful complement to lectures, slides, and practice problems. These quantitative trends align with open-ended reports describing the AICC as a course-aligned study aid that generates exam-style practice, offers step-by-step explanations, and helps organize lecture content for review.

Table 2: Likert Scale Results of Student Survey

Strongly agree | Somewhat agree | Neither agree nor disagree | Somewhat disagree | Strongly disagree

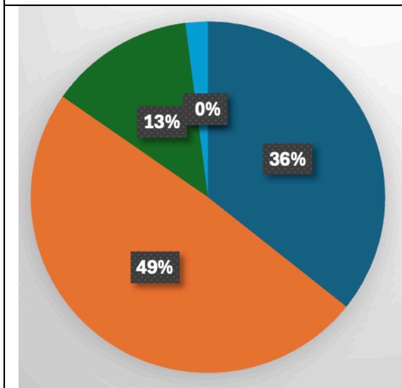

The AI Course Companion was easy to use.

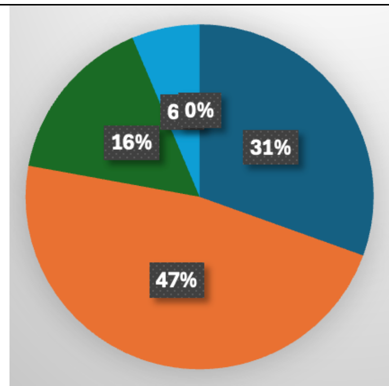

The tool's responses were clear and understandable.

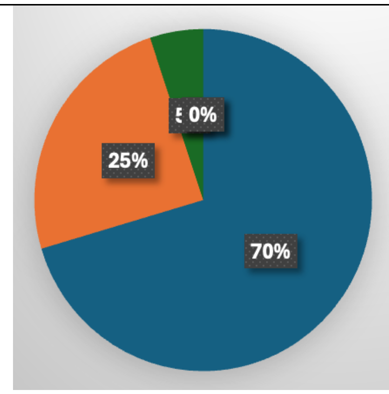

The AI responses were relevant to the course material.

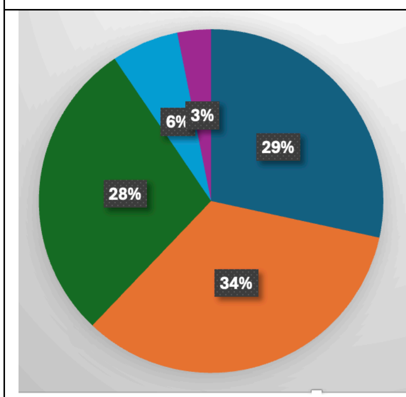

The AI tool saved me time when studying or reviewing course material.

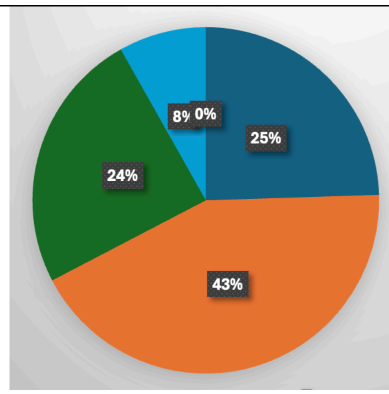

Using the AI Course Companion helped me better understand the course content.

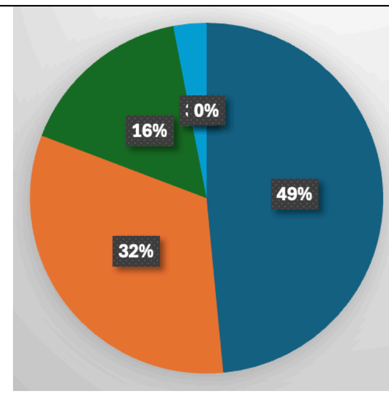

I would recommend this AI tool for future versions of the course.

Strongly agree | Somewhat agree | Neither agree nor disagree | Somewhat disagree | Strongly disagree

Finally, the global recommendation item provides a concise summary of overall sentiment. On the statement "I would recommend this AI tool for future versions of the course," 81% of surveyed students indicated agreement or strong agreement, while only a small minority reported that they would not recommend continued use. This pattern, reflected in the corresponding chart, suggests broad endorsement among active users and supports the interpretation that, despite noted areas for improvement, students perceived the AI Course Companion as a valuable addition to their courses.

4.3. Student survey – Open-Ended Answers Discussion

Students most commonly described the AI Course Companion as a course-aligned study aid that (a) generated practice opportunities (single problems, quizzes, and mock exams) in formats similar to graded work, (b) provided step-by-step explanations and answer checking that clarified procedures and repaired misconceptions, and (c) summarized and organized lecture materials to streamline exam preparation and navigation of course resources. Respondents emphasized that the tool's grounding in their slides and terminology distinguished it from general chatbots, reducing setup time and increasing perceived relevance; many also noted time savings and more focused study plans as a result. Collectively, these patterns suggest the AICC supported both conceptual understanding and procedural fluency in ways that mapped closely to course demands.

When asked about limitations or frustrations, students most often pointed to addressable issues: occasional inaccuracies or shallow coverage in advanced topics, uneven or late grounding to slide decks, multi-turn interaction hiccups (e.g., the tool forgetting context or prematurely changing tasks), readability/streaming ergonomics on mobile devices, and intermittent latency. A subset reported no problems, and many accepted assessment-integrity constraints (e.g., no direct solutions to graded work) as appropriate. These concerns primarily reflect reliability, usability, and workflow gaps rather than a rejection of the approach, and they outline clear priorities for iterative improvement. The issues that were reported by the users have been fixed in the later versions of the AI tool.

Students' suggestions for future versions reinforced those priorities: earlier and tighter ingestion of lecture materials with slide-level citations; a more compact, numbered presentation with controls to pause/scroll during generation; a unified chat entry with clearer practice workflows (submit → evaluate → hint → solution); visual/diagram support for course topics and optional image upload; faster and more resilient performance; and transparent provenance and onboarding ("what it can do" examples). These recommendations align the tool with authentic study practices and provide a concrete roadmap to enhance usability, trust, and learning effectiveness.

4.4. Instructor Survey Outcomes

All participating instructors completed the end-of-term survey, and their Likert-scale responses were consistently positive. All instructors indicated that they are comfortable using a tool like this in future semesters. Across the four items displayed in the charts, no instructor selected "somewhat disagree" or "strongly disagree," indicating that none perceived the AI Course Companion as misaligned or detrimental. Instead, ratings clustered at the "somewhat agree" and "strongly agree" categories, suggesting broad comfort with the tool and its role in their courses.

Items focused on alignment and integration showed the strongest consensus. All instructors agreed that the AI tool integrated well with their course materials and delivery, with responses split between "somewhat" and "strongly" agree. Similarly, instructors rated the AI-generated responses as accurate and aligned with course content, with most choosing agreement and a minority selecting a neutral option; no one indicated that the tool's answers were inaccurate or off-topic.

Perceptions of instructional value and workload impact were also favorable. On the item asking whether students appeared to benefit from the tool, based on observed questions, engagement, or feedback, most instructors agreed, and the remaining responses were neutral rather than negative. Instructors likewise agreed that the AI Course Companion reduced the number of repetitive questions or administrative tasks they typically receive, with all responses falling in the agree range. Together, these patterns indicate that instructors saw the AICC as both pedagogically aligned and practically helpful, supporting student learning while modestly easing routine instructional burden and reinforcing their willingness to continue using similar tools in future semesters.

Table 3: Likert Scale Results of Instructor Survey

Strongly agree | Somewhat agree | Neither agree nor disagree | Somewhat disagree | Strongly disagree

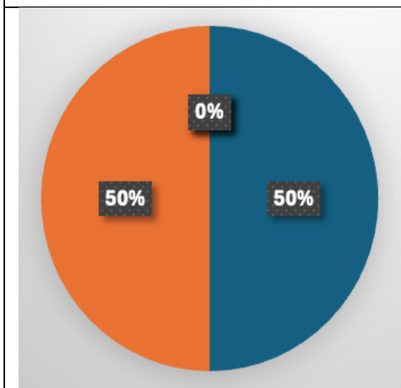

The AI tool integrated well with my course materials and delivery.

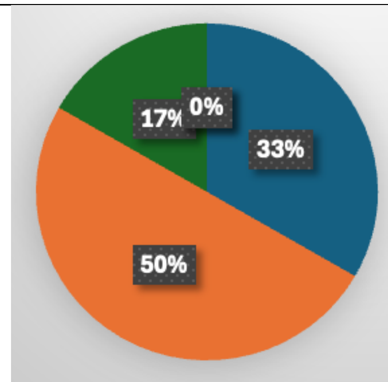

The responses provided by the AI tool were accurate and aligned with course content.

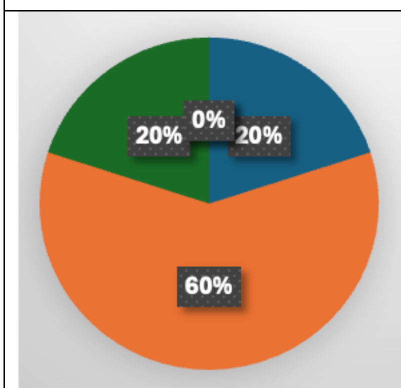

Students appeared to benefit from using the tool (based on questions, engagement, or feedback).

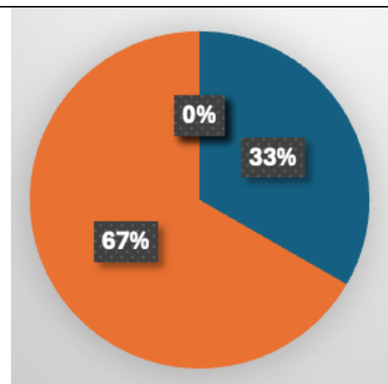

The tool reduced the number of repetitive questions or administrative tasks I typically receive.

Strongly agree | Somewhat agree | Neither agree nor disagree | Somewhat disagree | Strongly disagree

4.5. Instructor survey – Open Ended Answers Discussion

Based on open-ended responses from instructors, perceptions of the AI Course Companion were broadly positive. Instructors reported that students used the tool primarily for exam preparation and concept clarification, valuing its ability to generate practice questions, provide explanations (including multilingual support), and scaffold self-directed study. Several instructors noted that students leveraged the AICC to create additional study materials (e.g., practice problem sets), and no respondent reported negative student comments.

Instructors commonly characterized the AICC as a “24/7 teaching assistant,” citing real-time support that reduced repetitive questions, especially for foundational or technically demanding topics (e.g., stacks, assembly). The tool was viewed as particularly helpful in abstract or theory-intensive courses (e.g., Automata Theory), where immediate, personalized explanations supported engagement and comprehension. At the same time, instructors observed that a subset of students engaged less actively, underscoring the need for stronger motivation and onboarding.

Implementation was generally smooth with minimal logistical issues. One instructor experienced a temporary malfunction when uploading new materials, suggesting the value of automated notifications to instructors when ingestion errors occur. Instructors also indicated that effectiveness would be enhanced by tighter synchronization with weekly lectures, greater transparency about what materials are currently included, and more seamless alignment with evolving course content. Increasing student familiarity and confidence with the tool was identified as a complementary priority.

Overall, instructors reported that the AICC aligned with their pedagogical aims by reinforcing course-specific learning objectives, supporting asynchronous and online formats, and freeing class time for higher-order discussion and problem solving. They emphasized that the tool could guide responsible student use of AI, supporting conceptual understanding rather than providing rote answers, provided that clear course guidelines are established.

For future development, instructors recommended deeper integration with course content; automatic generation of practice exams by topic or slide selection; instructor-level controls to “tune” responses toward known areas of difficulty; stronger cross-document linking (slides, labs, readings); and an onboarding assistant to model effective prompting. While several respondents already described the design and usability as strong, they anticipated further gains with earlier (pre-semester) integration into course planning.

All instructors indicated they would recommend the AICC to colleagues, highlighting its potential to increase engagement, improve comprehension of complex topics, and reduce routine instructional load. Collectively, these qualitative findings suggest that the AI Course Companion is a scalable, course-aligned support for individualized learning in higher education.

4.6. Instructor Interviews

Semi-structured interviews with three instructors (Programming Concepts; Advanced Databases; Computer Organization/Logic Design) provided convergent evidence that the AI Course Companion functioned as a course-aligned study aid embedded in Canvas. Instructors typically introduced the tool via in-class announcements and brief demonstrations, highlighting its grounding in course materials; several uploaded review decks, past exams, or textbook chapters to prime course-specific responses. Reported student use clustered around assessment periods: one instructor observed rising use nearer exams, another estimated that ~30% of students engaged extensively while the remainder used it little, and a third suspected limited uptake in a small summer offering.

Perceived learning benefits centered on focused exam preparation, stepwise clarification of concepts, and reduced distraction relative to general-purpose chatbots because responses were constrained to approved course sources. Instructors noted confidence gains before exams and particular value for repeating students and true beginners in fundamentals, as well as for concept-heavy or theory-oriented topics. Where systematic performance comparisons were not feasible (e.g., small or accelerated summer cohorts), instructors nonetheless judged the AICC potentially useful when purposefully integrated with course study routines.

At the instructional level, the AICC did not require major changes to teaching approaches; however, instructors reported fewer repetitive questions and used the tool to extend practice opportunities between class meetings (e.g., creating review decks for AI-guided study). Overall pedagogical fit was described as strong when materials were current and well aligned with objectives.

Concerns were pragmatic rather than conceptual. Accuracy issues were minimal when the corpus was restricted to vetted materials, though stale uploads reduced the effectiveness of document-grounded chat. Adoption appeared sensitive to usability: one instructor attributed low uptake partly to a basic UI and suggested in-LMS nudges before exams. Others cautioned against over-reliance in principle but reported no incidents. Desired enhancements included stronger item-generation for quizzes/exams, instructor controls for exam-safe toggling, step-by-step tutoring that withholds final answers, and an updated, more intuitive interface with clearer linkage across slides, labs, and readings.

All three instructors indicated willingness to use or recommend the AICC going forward; immediately where alignment and materials are strong, and contingent on modest UX/feature improvements elsewhere, citing its potential to scaffold individualized practice, clarify difficult concepts, and reduce repetitive help-seeking outside class.

5. Discussion

5.1. Principal Findings

This mixed-methods pilot provides convergent evidence that a course-specific, LMS-embedded AI companion can be feasible and valued when deployed as an optional, ungraded resource in authentic undergraduate computing courses. Across student and instructor feedback, the overall pattern was strongly favorable, with endorsement for continued use paired with specific, actionable improvement needs. Among the recommendation items, 81% of student users and 100% of instructors supported offering the tool in future course iterations, suggesting broad acceptability among active users and strong instructor willingness to continue adoption.

Student results indicate that perceived value was driven by tight alignment to course artifacts and exam preparation workflows. Students described the AICC as most helpful when it generated exam-style practice opportunities, provided step-by-step explanations, and summarized or organized lecture materials for more efficient studying. Importantly, students explicitly differentiated this tool from general-purpose chatbots by emphasizing that grounding in their slides and course terminology reduced setup time and increased relevance, thereby supporting more focused study plans and perceived gains in both conceptual understanding and procedural fluency.

Instructor evidence corroborated these student-reported benefits and added a practical instructional-impact dimension. Instructors characterized the system as a “24/7 teaching assistant,” reporting fewer repetitive questions and observing that students used the tool primarily for concept clarification and exam preparation. Interviews further suggested that perceived effectiveness depended on up-to-date, well-aligned materials; when course content was current and grounded, accuracy concerns were limited, whereas stale or incomplete uploads reduced the benefits of document-grounded support.

Finally, both interview and analytic summaries point to a distinctive temporal use pattern: engagement increased near assessment periods, with usage analytics (and instructor observation) indicating spikes in practice-question activity in the days preceding exams. This aligns with the tool's design emphasis on retrieval practice and exam-style preparation, suggesting that students treated the AICC as a targeted study aid rather than a novelty feature.

5.2. Interpretation and Implications

A core interpretation of these findings is that "course specificity" is not merely an implementation detail, but a primary mechanism of perceived usefulness and trust. Students and instructors repeatedly attributed value to the system's tight alignment with course artifacts (slides, homework, exam style) and its ability to keep help-seeking "on-topic." From a learning-sciences perspective, this alignment plausibly reduces extraneous cognitive load (less effort translating generic explanations into course conventions), while enabling more direct engagement with germane processing (connecting explanations to the exact representations used in class). This interpretation is consistent with the system's intended design choices: lecture grounding, worked-example style explanations, and retrieval-practice-oriented item generation.

A second implication concerns how AI support can be structured to reinforce, rather than replace, productive student thinking. Students valued stepwise explanations and answer checking, while also acknowledging (and often accepting) integrity-related constraints such as not providing direct solutions to graded work. This suggests that course-specific AI companions may be best positioned as formative tools that (a) offer scaffolded reasoning, (b) provide hints and intermediate steps, and (c) encourage reflection or self-explanation before revealing final answers. In practice, this implies designing "exam-safe" and "learning-first" interaction modes (e.g., hint-first tutoring; rubric-aligned feedback on a student's attempt; guided solution outlines) that are explicitly communicated to students during onboarding.

Design implications from both student and instructor feedback converge on four priorities. First, deepen and make visible the grounding process: earlier ingestion of course materials, slide-level citations, and transparency about which documents are currently included were recurring needs. Second, strengthen structured practice workflows: students asked for a unified practice loop (submit → evaluate → hint → solution), and instructors requested auto-generated practice exams by topic or slide selection. Third, address usability and access friction: mobile readability, controls to pause/scroll during response generation, and latency/resilience were salient. Fourth, implement "adoption scaffolds" for both roles: a student onboarding assistant to model effective prompting and a planned pre-semester integration workflow for instructors (so materials and expectations are ready from the start).

For instructors and administrators, the results support several pragmatic recommendations for responsible integration. Because the system was optional and ungraded, uptake was likely shaped by students' study habits and the clarity of instructor messaging; interviews suggested that proactive students and those preparing for exams engaged most. Accordingly, implementation should include: (1) a brief in-class demonstration early in the term; (2) explicit syllabus language clarifying appropriate use (e.g., concept clarification, practice generation, exam preparation) and boundaries (e.g., graded-work restrictions); (3) periodic, low-friction "nudges" tied to authentic course rhythms (e.g., pre-exam review weeks); and (4) equity and accessibility checks (mobile usability, disability accommodations, and clear guidance for students with lower AI literacy). While the present study cannot establish causal learning gains, the endorsement and use patterns suggest that course-specific AI companions can be a credible component of a broader student-success strategy when paired with clear pedagogical framing and transparent grounding.

5.3. Strengths and Limitations

This study has several strengths. First, it evaluated the AI Course Companion in authentic conditions across six undergraduate computing courses, providing ecological validity relative to studies limited to single-course pilots. Second, the mixed-methods approach enabled triangulation across student survey results, instructor survey responses, and instructor interviews, strengthening confidence in the convergent themes around alignment, exam preparation, and reduced repetitive questions. Third, system logging was anonymous by design, which supports privacy-preserving analytics in real deployments.

Key limitations should be considered when interpreting the findings. The evaluation relied heavily on self-reported perceptions, and the analytic inclusion focused on students who used the tool at least occasionally (three or more times), creating risks of self-selection and response-bias. The single-institution context and single-term deployment constrain generalizability, particularly given disciplinary specificity (computing) and the accelerated nature of a summer semester. Additionally, while usage logs captured feature-level engagement, the anonymization approach precluded linking usage to individual learning outcomes (e.g., grades, assignment performance, or persistence), limiting the study's ability to estimate causal impacts. Finally, the design was non-experimental (no randomization or comparison group), so findings should be interpreted as evidence of feasibility, acceptability, and perceived value rather than proof of efficacy.

5.4. Future Work

Future work should prioritize iterative redesign aligned with the improvement roadmap surfaced by students and instructors: (a) tighter and more transparent lecture grounding (including slide-level citations and real-time corpus status), (b) structured practice and "practice exam by topic/slide" generation, (c) improved cross-document linking across slides, labs, and readings, and (d) stronger onboarding supports that model effective prompting and appropriate use.

Methodologically, the next step is to move from feasibility/acceptability toward outcome-focused designs. With appropriate consent and privacy safeguards, future studies could connect usage intensity and feature pathways (e.g., practice generation vs. concept Q&A) to course outcomes (exam performance, assignment success, withdrawal rates), enabling stronger inferences about which affordances matter most for which students. Complementary experimental or quasi-experimental designs (e.g., section-level rollouts, staggered introductions, encouragement designs) could help disentangle selection effects from tool effects while preserving authentic deployment conditions.

Finally, cross-course and cross-institution replication is essential. While this pilot spanned multiple computing courses and modalities (including an online asynchronous course), broader evidence is needed across disciplines, instructor styles, and institutional contexts. Such work should also examine equity considerations (who adopts, who benefits, and under what supports), as well as governance questions (instructor controls, integrity modes, and institutional policies) that will determine whether LMS-embedded AI companions can scale responsibly.

6. Conclusion

This mixed-methods pilot suggests that an LMS-embedded, lecture-grounded AI Course Companion can be feasibly deployed as an optional, ungraded resource in authentic undergraduate computing courses and is perceived as valuable by both students and instructors. Across six courses (311 enrolled students) with survey feedback from 62 students who used the tool at least three times and participation from all five instructors, endorsement for continued use was high (81% of students; 100% of instructors). Students emphasized benefits tied to course alignment; exam-style practice generation, stepwise explanations, and targeted summaries grounded in their materials; while instructors characterized the tool as a "24/7 TA" that supported individualized practice and reduced repetitive help-seeking when course artifacts were current and well aligned.

At the design level, the findings reinforce that course specificity and grounding to vetted instructional artifacts are central to perceived relevance, usefulness, and trust. Usage patterns (notably increased practice-question engagement prior to exams) suggest that students primarily adopted the AICC as a targeted study aid aligned with assessment preparation. Together, these results motivate continued development and evaluation of course-specific AI companions, with near-term priorities including deeper and more transparent lecture integration, stronger practice-exam generation by topic/slide, improved linking across slides/labs/readings, and onboarding supports that encourage learning-first, integrity-aware use.

Acknowledgments

We thank the Bellini College of AI, Cybersecurity, and Computing at the University of South Florida for its support and partnership in developing and deploying the AI Course Companion. We especially thank Dr. Sudeep Sarkar, Launch Dean of the Bellini College of AI, Cybersecurity, and Computing, for his continued support of the project. We also acknowledge the contributions of volunteer students Juan Franco Boeta, Zyad Abd-Elrahman, and Akash Patel, and project developers Harshavardan Yuvaraj and Mahek Patel, for implementing key AI Course Companion features.

Conflicts of Interest

The authors declare no conflicts of interest.

References

[1] M. Amoozadeh et al., "Trust in Generative AI among Students", Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE 2024), March 20–23, 2024, Portland, OR, USA. https://doi.org/10.1145/3626252.3630842

[2] L. Labadze, M. Grigolia, and L. Machaidze, “Role of AI chatbots in education: systematic literature review,” *Int. J. Educ. Technol. High. Educ.*, vol. 20, Art. no. 56, 2023, doi: 10.1186/s41239-023-00426-1.

[3] N. Maleki, B. Padmanabhan, and K. Dutta, “AI hallucinations: a misnomer worth clarifying,” in *Proc. 2024 IEEE Conf. on Artificial Intelligence (CAI)*, 2024, pp. 133–138, doi: 10.1109/CAI59869.2024.00033.

[4] O. O. Ayeni, N. M. Al Hamad, O. N. Chisom, B. Osawaru, and O. E. Adewusi, “AI in education: a review of personalized learning and educational technology,” *GSC Adv. Res. Rev.*, vol. 18, no. 2, pp. 261–271, 2024, doi: 10.30574/gscarr.2024.18.2.0062.

[5] R. Liu, C. Zenke, C. Liu, A. Holmes, P. Thornton, and D. J. Malan, “Teaching CS50 with AI: leveraging generative artificial intelligence in computer science education,” in *Proc. 55th ACM Tech. Symp. Comput. Sci. Educ. (SIGCSE ’24)*, 2024, pp. 750–756, doi: 10.1145/3626252.3630938.

[6] S. Kahl *et al.*, “Evaluating the impact of advanced LLM techniques on AI-lecture tutors for a robotics course,” *arXiv* preprint arXiv:2408.04645, 2024.

[7] W. Lyu, Y. Wang, T. (R.) Chung, Y. Sun, and Y. Zhang, “Evaluating the effectiveness of LLMs in introductory computer science education: a semester-long field study,” in *Proc. 11th ACM Conf. Learning @ Scale (L@S ’24)*, 2024, pp. 63–74, doi: 10.1145/3657604.3662036.

[8] D. R. Thomas *et al*., “Improving student learning with hybrid human-AI tutoring: a three-study quasi-experimental investigation,” in *Proc. 14th Learning Analytics and Knowledge Conf. (LAK ’24)*, 2024, pp. 404–415, doi: 10.1145/3636555.3636896.

[9] P. Shah, *AI and the Future of Education: Teaching in the Age of Artificial Intelligence*. Hoboken, NJ, USA: John Wiley & Sons, 2023.

[10] J. Sheard et al., “Instructor Perceptions of AI Code Generation Tools – A Multi-Institutional Interview Study”, Proceedings of the 55th ACM Technical Symposium on Computer Science Education (SIGCSE 2024), March 20–23, 2024, Portland, OR, USA. https://doi.org/10.1145/3626252.3630880

[11] H. Soliman, M. Kravcik, A. T. Neumann, Y. Yin, N. Pengel, and M. Haag, “Scalable mentoring support with a large language model chatbot,” in *Proc. 19th Eur. Conf. Technol. Enhanced Learn. (EC-TEL 2024)*, vol. 15160, *Lect. Notes Comput. Sci.* Cham, Switzerland: Springer, 2024, pp. 260–266, doi: 10.1007/978-3-031-72312-4_37.

[12] H. Wang *et al.*, “Examining the applications of intelligent tutoring systems in real educational contexts: a systematic literature review from the social experiment perspective,” *Educ. Inf. Technol.*, vol. 28, no. 7, pp. 9113–9148, 2023, doi: 10.1007/s10639-022-11555-x.