Alan: A Classroom Deployment of An Educational Chatbot That Cites Resources

Vaibhav Seshadri
v.seshadri@mail.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Yassine Berrachid
yassine.berrachid@mail.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Ethan Hugh
ethan.hugh@mail.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Abhishek Madan
ab.madan@mail.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Salma Emara
salma@ece.utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Abstract

The rise of large language models (LLMs) such as ChatGPT presents new opportunities and challenges for computing education. While these models can provide instant assistance, their generality and lack of source grounding can produce inaccurate or unverifiable responses. This extended abstract introduces Alan, a course-specific virtual teaching assistant built on a retrieval-augmented generation (RAG) framework. Alan answers student questions using verified course resources such as lecture slides, lab handouts, notes, and recordings, and cites these sources directly in each response. Deployed in a large introductory programming course, Alan was used by over 250 students and logged more than 2,000 queries. A mixed-methods analysis combining usage data and survey responses shows that students found Alan most helpful for conceptual and administrative questions, valuing its transparency and course alignment. We conclude with lessons learned about integrating grounded AI support in computing classrooms.

CCS Concepts

• **Applied computing** → **E-learning**; • **Computing methodologies** → *Natural language processing*; • **Human-centered computing** → **Human computer interaction (HCI)**.

Keywords

AI in Education, Educational Technology, Retrieval-Augmented Generation (RAG), Educational Chatbots, Large Language Models (LLMs), Human-AI Collaboration

ACM Reference Format:
Vaibhav Seshadri, Yassine Berrachid, Ethan Hugh, Abhishek Madan, and Salma Emara. 2026. Alan: A Classroom Deployment of An Educational Chatbot That Cites Resources. In *Proceedings of the 57th ACM Technical Symposium on Computer Science Education V.2 (SIGCSE TS 2026), February 18–21, 2026, St. Louis, MO, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3770761.3777344

SIGCSE TS 2026, St. Louis, MO, USA

ACM ISBN 979-8-4007-2255-4/2026/02
https://doi.org/10.1145/3770761.3777344

1 Introduction

Generative AI tools have rapidly entered educational contexts, offering instant feedback and natural-language explanations. Researchers at Harvard and Yale University [3] and others [1, 2, 4] increasingly explore using educational chatbots as virtual teaching assistants, especially in high-enrollment courses where student to teacher ratio is high and timely responses or personalized feedback is not attainable. Yet, their lack of verifiable grounding has raised legitimate concerns: instructors worry that students might treat unverified AI output as authoritative, while students struggle to judge when a response can be trusted. To address these issues, we developed Alan, a course-specific chatbot that generates answers solely from instructor-approved materials and displays clickable citations to those resources. Alan is designed to combine the responsiveness of general LLMs with the reliability and transparency required for academic settings.

In Winter 2025, we deployed Alan in a first-year programming course (~450 students) to examine how students engage with a verifiable AI assistant and how they perceive its role alongside existing help channels such as discussion forums.

Our contribution is threefold:

- We describe a scalable design for a retrieval-augmented chatbot tailored to course material.
- We report patterns of student usage and perceived helpfulness from a real-world classroom deployment.
- We discuss design implications for trustworthy AI integration in computing education.

2 System Design

Design Goals. Alan was designed to (1) ensure every answer is based on course content, (2) support instructor oversight, and (3) build student trust through transparency. The system therefore retrieves relevant course resources, generates a response using a large language model, and cites those resources inline and at the end of each answer. The instructor can observe student queries and answers to ensure they are answered correctly.

Architecture. Alan follows a retrieval-augmented generation (RAG) workflow with three stages. **Retrieval.** The system searches a hybrid index of course materials using both keyword and semantic similarity methods. **Augmentation.** Top-ranked excerpts are combined with the student's query to form a context-rich prompt.

Generation. A GPT-4-class model produces an answer, automatically inserting citations to the retrieved excerpts.

The knowledge base integrates slides, lecture transcripts, lab instructions, websites, and LMS files, updated automatically throughout the term.

Interface and Features. Alan is a web-based chatbot accessible on desktop and mobile. Students log in, join their course using an instructor-provided code, and may opt-in to research data logging. Each answer displays clickable in-text citations and a list of source documents to promote verifiability.

Students can rate responses, request regeneration, or use a Search Mode to explore relevant course materials directly. Instructors can view anonymized dashboards summarizing active users and frequently asked questions to monitor learning trends.

3 Classroom Deployment Findings

Deployment. Alan was deployed in a Winter 2025 offering of an introductory C-programming course. Out of 453 students, 251 registered for Alan and 150+ consented to logging. From late January to April, Alan recorded 2,217 queries with anonymized user IDs and timestamps. After the final exam, a voluntary survey (40 responses) captured student perceptions of helpfulness and comparison with other GenAI tools.

User Data. Student engagement with Alan mirrored the course rhythm, spiking around weekly lab deadlines, the midterm, and the final exam. Query activity peaked both around noon and again between 10 p.m. to midnight hours when human help was less accessible, which indicates that Alan's role as an after-hours study companion.

While most students (100) asked only a few questions, a smaller group (10) accounted for 50 or more queries each, forming a power-user cohort who relied on the chatbot for continuous support.

Using qualitative categorization, student queries clustered into six groups: Conceptual (63%), e.g., "How do I count digits in a number?" Lab (17%), e.g., "How do I approach Lab 5?" Administrative/logistic (7%), e.g., "When are office hours?" Lectures (5%), Past exams (%), and Other (4%). Lab handouts were among the most frequently retrieved documents, underscoring that conceptual and lab help were Alan's dominant use cases.

Fig. 1 shows the number of queries sent to Alan and posts on Piazza, the course discussion forum, throughout the term. Both systems show clear spikes in activity around due dates. Piazza remained active throughout the term. The overall number of forum posts per student was comparable to the previous year without Alan, suggesting that the chatbot supplemented rather than replaced human support. Students appeared to turn to Alan for quick clarification and Piazza for complex or social help-seeking.

Perceived Helpfulness. Survey results reflected generally positive perceptions. On a 5-point Likert scale, mean helpfulness scores were 4.2 for conceptual questions, 3.7 for problem-solving, 3.9 for administrative questions, and 4.0 overall.

Students highlighted three valued features: **Course-specific grounding.** "I enjoyed that it provided sources to course materials, especially timestamps to lecture recordings," **Trust and accuracy.** "I am sure Alan has the most accurate definitions of concepts learned in this course," **Immediacy.** "It helped to have Alan as the first

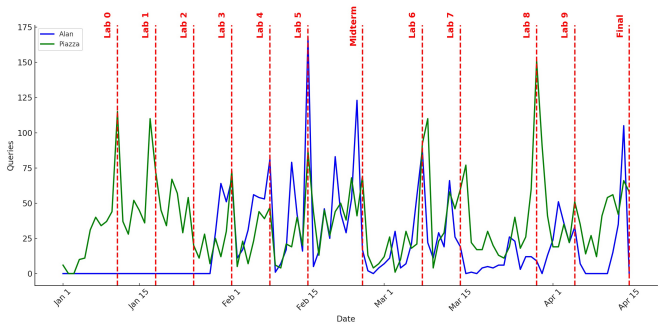

Figure 1: Number of queries to Alan and posts on Piazza throughout the term, with spikes around lab deadlines, the midterm, and the final exam. The number of queries to Alan was 0 before January 29 as no data was tracked.

resource before asking questions on Piazza." These responses indicate that transparency and context alignment fostered trust and positioned Alan as a reliable first-line resource.

Limitations, Lessons and Future Work. Students noted weaker performance on open-ended or visual questions, such as those requiring interpretation of hand-drawn diagrams or guidance for approaching lab tasks. These limitations stem from restricted access to image data and guardrails preventing out-of-context reasoning.

Future work will explore multimodal retrieval (e.g., diagrams, handwritten notes), longitudinal impact across multiple courses, and instructor analytics to identify emerging student misconceptions in real time.

4 Concluding Remarks

This work presents Alan, a retrieval-augmented, course-specific educational chatbot that cites its sources. In a large classroom deployment, students valued its immediacy, accuracy, and alignment with course materials. Although limitations remain in multimodal reasoning and complex lab support, our findings demonstrate the promise of grounded AI systems to enhance learning without displacing human interaction. By emphasizing transparency and trust, Alan offers a replicable model for integrating generative AI responsibly into computing education.

References

[1] Arto Hellas, Juho Leinonen, and Leo Leppänen. 2024. Experiences from integrating large language model chatbots into the classroom. In *Proceedings of the 2024 ACM Virtual Global Computing Education Conference V. 1*. 46–52.

[2] David Lau and Michael Cooper-Stachowsky. 2024. Architecture and usage information for a course-aligned, GPT-powered chatbot. In *Proceedings of the Canadian Engineering Education Association (CEEA)*.

[3] Rongxin Liu, Carter Zenke, Charlie Liu, Andrew Holmes, Patrick Thornton, and David J. Malan. 2024. Teaching CS50 with AI: Leveraging Generative Artificial Intelligence in Computer Science Education. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education (SIGCSE 2024), Vol. 1*. ACM, Portland, OR, USA, 1–7. doi:10.1145/3626252.3630938

[4] Mike Wu, Noah D. Goodman, Chris Piech, and Chelsea Finn. 2021. ProtoTransformer: A Meta-Learning Approach to Providing Student Feedback. *arXiv preprint arXiv:2107.14035* (2021). https://arxiv.org/abs/2107.14035