Paper ID #51355

Improving Academic Student Success: Using AI Chatbots as Virtual Teaching Assistants in Petroleum Engineering

Dr. Jason M. Keith, Iowa State University of Science and Technology

Jason Keith is Senior Vice President and Provost at Iowa State University. Prior to this appointment, he was Dean of the Bagley College of Engineering at Mississippi State University. Keith was recognized in 2014 as a Fellow of ASEE.

HIMANSHU TRIPATHI
Charlotte Elise Crowell, The University of Alabama
Shahram Rahimi, Mississippi State University

Improving Academic Student Success: Using AI Chatbots as Virtual Teaching Assistants in Petroleum Engineering

Abstract

The landscape of higher education is already being reshaped by Artificial Intelligence (AI) and Large Language Models (LLMs), and university operations are being impacted in many ways. One particular aspect that has showed promise in colleges of engineering and computer science is the use of AI and LLM to help students in departmental entry-level courses, many of which have historically been difficult for students to successfully navigate. In the past these courses were viewed as gateway courses, without which a student would not stay in the major. However, over the past couple of decades, more emphasis has been placed on the scholarship of teaching and learning, supplemental instruction, experiential learning and the development of novel methods to help improve overall academic student success. While all of these strategies are helpful, we aim to supplement these positive impacts via our work in the development of AI and LLM to improve student success in a key course taken by petroleum engineering (and often by chemical engineering) students. Offered as an interactive resource to struggling students, to reaffirm performance of successful students, or serve as a study aid, the bot guides students through problem identification, developing solution methods, working carefully through key steps and verifying solutions. Instructors can use these tools to run virtual help sessions, exam study hours and serve as a reference / resource material to review content that was taught earlier in the semester and may need to be recalled by students when solving more complex design problems in the field.

Introduction and Literature Review

Artificial Intelligence (AI) has emerged as a transformative force in higher education, most notably within engineering and computing disciplines worldwide. Beyond its prominence and utility in research areas such as autonomous systems in manufacturing and transportation, materials science, energy systems and big data, AI also offers significant opportunities to enhance teaching and learning. Chatbots used in higher education present unique and compelling advantages: they provide round-the-clock support, allow students to ask anonymous questions, especially valuable at reducing student hesitation to speak up or ask what may be considered a “bad” question [1-4]. Furthermore, chatbots can also effectively deliver interactive tools like quizzes to reinforce knowledge retention. Finally, these systems can manage large volumes of inquiries efficiently, alleviating the burden on faculty who often face time constraints in responding to student needs, all while adding a degree of personalization, and are especially valuable when limited resources prohibit the broad availability of teaching assistants.

Early applications of AI in education demonstrated its potential to support teaching and learning through conversational agents. Georgia Tech’s Jill Watson [5] and subsequent work at Stanford [6] pioneered the use of AI-driven question and answer (Q&A) systems to assist students, showing strong initial performance and suggesting long-term promise. Building on these initial foundations, chatbots have been used from static Q&A systems to become more interactive and personalized tools aimed at enhancing efficiency and student engagement [7-9]. Specific examples include proprietary AI services such as Khan Academy’s Khanmigo [10] and the University of Sydney’s

Cogniti [11], which provide tutoring and Q&A support across a variety of subjects such as chemistry, physics, and mathematics, and even advanced topics like anesthesiology and genetics. Broader literature underscores the educational benefits of chatbots, including improved accessibility, interaction, and retention of knowledge [12-17] including cost-effectiveness, responsiveness, and student preference for anonymity. These chatbots have shown that they can assist in improving student performance, and Alsafari's recent work [18] suggests that LLM algorithms are now developed enough to handle sophisticated student queries.

It should not be a surprise that universities are increasingly leveraging conversational agents to enhance student support and engagement [19-21]. Our prior work [22-24] has advanced initiatives to strengthen student success, beginning with the development of a chatbot to support recruitment and guide undecided students toward majors that match their interests and competencies. This system analyzed responses to recommend academic pathways and also connected students to Student Affairs resources, to help instill a feeling of belonging for every student. Building on this foundation, we shifted our focus from recruitment to persistence within high-risk courses. Specifically, efforts targeted an upper-level petroleum reservoir rock properties course and an introductory computer science course. In this work, we focus on the introductory petroleum engineering course, with the intent on improving interaction with students and increasing first-to-second year retention in the major. Specifically, we apply the use of basic engineering principles, including unit conversions, dimensional analysis, and empirical correlations to properly solve a multi-step problem. This tool is meant to tutor students through interactive multiple-choice quizzes, with an intended audience of students who want to study for their midterm or final exam or receive extra support if they struggle to obtain correct solutions in prior assignments.

We note studies have shown that frequent low-risk quizzes and well-designed multiple-choice question (MCQ) assessments can also significantly improve learning outcomes. When constructed thoughtfully, multiple-choice questions support higher order thinking by aligning with learning objectives [25-26]. Backward design and authentic assessment strategies have also been used to align goals and evaluation methods, as highlighted by Vanderbilt's resources on effective teaching practices [27]. Sotola and Credé [28] illustrate that frequent quizzes worth a low portion of the overall course grade are moderately associated with improved course performance and significantly increase the odds of passing, while Oc and Hassen [29] examine how multiple-answer questions promote deeper thought processing but are perceived as more challenging than single-answer formats.

Current State of the Algorithm

With each research / educational problem or application that we have approached, we have added functionality to our AI algorithm. The academic oriented bots [22-24] also utilize features from our work in AI bots applied to healthcare [30] and cybersecurity [31], and we will briefly outline the present state.

This study employs a Teaching Assistant (TA) chatbot designed to deliver real-time, adaptive support for students. The system addresses both conceptual questions—such as applying engineering principles—and practical challenges, including numerical tasks. The architecture follows a modular design consisting of interconnected components that classify and process queries efficiently. Core elements include a Query Router, Retriever, Concept Agent, MCQ Generator, Explanation Agent, Code Generator, Answer Verifier, and Feedback Mechanism, all integrated to work seamlessly with LLMs to generate accurate and context-aware responses.

Figure 1 illustrates the functionality of each key component of the chatbot architecture. Components outlined in green highlight new additions to the architecture originally presented in [22]. Labels T1-T3 show the components of the conceptual query resolver. T1 represents relevant course materials that the system may refer to when answering queries. T2 forms the retrieval pipeline that stores and queries the vectorized course materials. T3 is the LLM "brain" of the system, which can draw upon retrieved materials when answering theoretical questions. Labels N1-N6 highlight the components of the numerical query resolver. N1 extracts relevant information and subproblems from the given query, while N2 generates steps to solve each subproblem. N3 creates an MCQ for each subproblem to test the user's understanding, while N4 provides custom feedback to the user based on their response. N5 creates a code to numerically solve the given subproblem and N6 uses the output of that code to verify the user's solution.

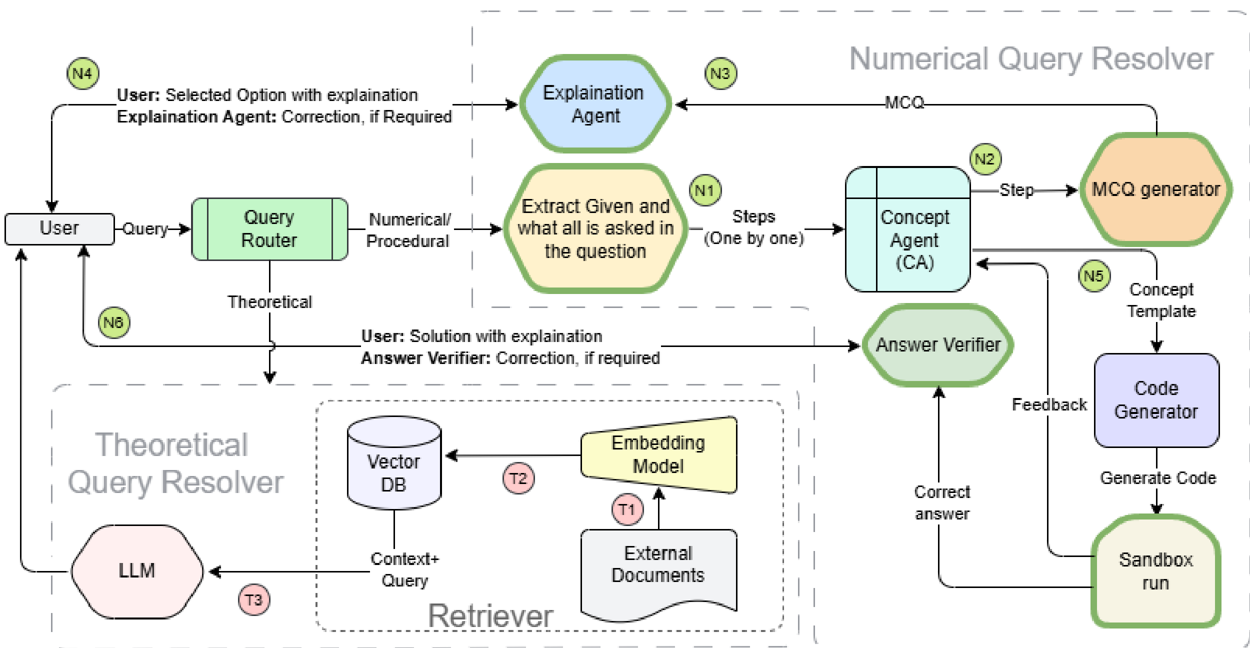

Figure 1: Architecture for the teaching assistant chatbot

The Query Router (Figure 2) classifies the user's query into either a theoretical or numerical/code-based query. This uses a lightweight classifier (a small fine-tuned LLM or by simple zero-shot prompting, which is a method that generates a response without any examples included within the prompt) to predict in a zero-shot approach, the query type based on its structure and keywords [22].

```
prompt = f"""

   Instruction: You are an AI assistant trained to classify user queries into two categories: Theoretical or Numerical-Procedural. Follow these guidelines:

   1. Theoretical: The query involves abstract concepts, principles, or qualitative reasoning. It seeks explanations, definitions, or discussions without requiring calculations, numerical results, or step-by-step procedures.

   2. Numerical-Procedural: The query involves numbers, calculations, data analysis, code, or step-by-step instructions. It seeks specific numerical answers, solutions to equations, code implementations, or procedural guidance.

   Task: Read the user query below and classify it as either Theoretical or Numerical-Procedural. Provide only the classification as your output.

   User Query: {query}

   Output:

   """
```

Figure 2: Zero-shot prompt for query classification [22]

For theoretical queries, the system applies a Retrieval-Augmented Generation (RAG) workflow (a connection to accessible data, rather than using a pre-trained model that uploads a document and retrieves relevant passages to support the LLM's response) to deliver accurate, context-aware responses. This process begins by retrieving relevant content from curated knowledge sources such as course materials and textbooks. The retrieved documents, combined with the user query, are then processed by an LLM, which synthesizes the information into a coherent and scientifically precise explanation.

After the Query Router classifies the user input as a numerical query, the numerical workflow begins by using a lightweight Small Language Model (SLM) to extract all given values, assumptions (Figure 3), and required intermediate targets from the user query. In this figure, Phi-4 extracts givens and subquestions, Phi-3 generates MCQ distractors and solves in a sandbox, OSS20B verifies choices, checks answers, gives feedback, and summarizes. In this figure, green darkens with the size of the sub-model.

Following the SLM step, our chatbot decomposes the problem into a sequence of sub-questions that collectively lead to the final objective (Figures 4 and 5). Note, that user entries are a reddish orange and the LLM output is in blue. Next, for each sub-question, the Concept Agent generates the correct step-by-step methodology required to solve that sub-question, while the MCQ Generator produces three plausible but incorrect methodologies and combines them with the correct one into a four-option MCQ. In future versions, in order to make this more accessible to students, teaching assistants and faculty, we will interface with other AI tools to digitize text from images. We also plan on incorporating a module on prompt engineering to best prepare students to effectively use this chatbot. Furthermore, as this tool evolves, we intend on building functionality to handle poorly designed prompts / inconsistent systems of equations.

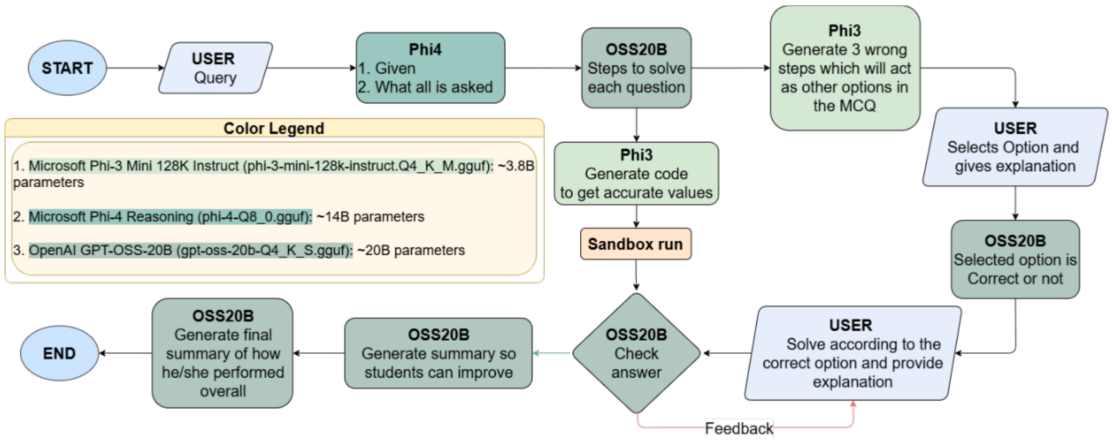

Figure 3: Workflow for guided numerical problem solving

> user_query = A pipeline transports gasoline at 300,000 barrels per day. The pipeline has an outside diameter of 24 inches and a wall thickness of 0.5 inches. Gasoline properties are: specific gravity = 0.74 and dynamic viscosity = 0.6 cP. Assume steady flow and use SI units where needed. Determine the Reynolds number for the flow in the pipeline by doing the following: convert the flow rate from barrels/day to gallons/min (1 bbl = 42 gal; 1 day = 1440 min), convert the flow rate to m^3/s (1 gal = 0.00378541 m^3; 1 min = 60 s), compute the inside diameter of the pipe and convert it to meters, compute the cross-sectional area of the pipe (m^2), compute the average velocity in the pipe (m/s), compute the fluid density from specific gravity ($\rho = SG \times 1000\ kg/m^3$), convert viscosity to Pa·s (1 cP = 0.001 Pa·s), and compute the Reynolds number ($Re = \rho VD/\mu$).

Figure 4: Numerical question posed by user

After the user selects the correct answer for the MCQ, the code generator agent generates a code to solve the subproblem using the selected steps. The input to the code generator is a concept template, which serves as a structured framework for problem-solving. Based on this template, the system produces Python code using only basic language constructs such as control flow, arithmetic operations, and standard library functions, avoiding complex dependencies to ensure simplicity and portability. For numerical problems, the generator applies the given information extracted by the SLM at the first step, maintaining clarity and correctness throughout the computation. The output is self-contained Python code capable of solving the specified numerical or programming subproblem. A simple example of this code generation is shown in our prior work [22].

The system then executes the generated code in a sandbox environment and uses the Answer Verifier to compare the sandbox-derived result to the user's computed answer for the same subproblem. If the values match within tolerance, the system proceeds to the next subproblem; otherwise, the Answer Verifier triggers the Feedback Mechanism, which provides corrective hints that help the user identify the mistake (e.g., inconsistent units, incorrect intermediate quantity, or formula misuse) while still encouraging the user to recompute rather than just revealing the final answer. This feedback loop repeats until the user arrives at the correct solution for the subproblem, at which point the workflow advances to the next subproblem in the sequence. We detail an example corresponding to the same query and concept template in the next section.

GIVEN:
- The pipeline transports gasoline at 300,000 barrels per day.
- The pipeline has an outside diameter of 24 inches.
- The pipeline has a wall thickness of 0.5 inches.
- Gasoline specific gravity = 0.74.
- Gasoline dynamic viscosity = 0.6 cP.
- Assume steady flow.
- Conversion factors:
- 1 barrel = 42 gallons.
- 1 day = 1440 minutes.
- 1 gallon = 0.00378541 cubic meters.
- 1 minute = 60 seconds.
- 1 centipoise (cP) = 0.001 Pascal-seconds (Pa·s).
- Use SI units where needed.

Sub-Questions:
1. the flow rate from barrels per day to gallons per minute: (300,000 barrels/day) × (42 gallons/barrel) / (1440 minutes/day)
2. the Reynolds number: Re = ρVD/μ (use ρ = SG×1000, μ from cP→Pa·s, D = inside diameter in meters, V = Q/A with Q in m^3/s))

Figure 5: Extracted given and what all sub-questions asked in question illustrated in Figure 4

Example and Discussion

In Figure 6, we show the correct flow for one of the subproblems related to this question. The user works out the correct solution and is encouraged to move on to the next subproblem. In Figure 7, we show how the student has chosen an incorrect solution and is guided by the system to recompute and move towards the correct answer.

We want to add additional remarks regarding the implementation details, where we utilized one dataset for this research: the course curricula of "Introduction to Petroleum Engineering." All development was implemented using the LangChain framework, enabling seamless integration of data preprocessing, embedding, retrieval, and response generation. Course materials, provided as PDF files, were preprocessed using LangChain's DirectoryLoader and PyPDFLoader. Documents were split into manageable chunks using the Recursive Character Text Splitter with a chunk size of 1000 characters and an overlap of 200 characters to preserve contextual integrity. These chunks were embedded using OpenAI's *text-embedding-ada-002* model, and the resulting vectors were stored in Chroma DB, a high-performance vector database.

SUB-QUESTION 1: the flow rate from barrels per day to gallons per minute: (300,000 barrels/day) × (42 gallons/barrel) / (1440 minutes/day)

How would you solve this? Choose the best METHOD option (A/B/C/D):
A) Add 42 to 300,000 and then divide by minutes per day to get gallons/min.
B) Multiply 300,000 by 42 to get gallons/day, then divide by 1440 to get gallons/min.
C) Divide 300,000 by 42 to get gallons/day, then multiply by 1440 to get gallons/min.
D) Multiply 300,000 by 1440 to get barrels/min, then multiply by 42 to get gallons/min.

Your choice (A/B/C/D): B
Why do you think this method is correct? (bbl/day)*(gal/bbl)/(min/day)=gal/min so multiply by 42 then divide by 1440.

Tutor feedback:
Correct! Multiply by gallons-per-barrel, then divide by minutes-per-day.

Now solve it.
Your final answer (number; you may include units): 8750 gal/min
Your work/explanation (optional): 300000*42/1440 = 8750 gal/min

Tutor feedback:
Correct! Nice unit cancellation and arithmetic.

Summary:
Great job! 10/10 – clear, correct conversion.

Figure 6: Sub-question 1, interacting with the student / user who solves the question correctly

For retrieval, we employed LangChain's vector store abstraction with similarity search configured to retrieve contextually relevant documents. OpenAI's GPT-4 API, integrated via LangChain, was used to generate accurate, context-aware responses. We note that we also experimented with open source LLMs like GPT-OSS-20B, in order to ensure cost-effectiveness. This pipeline maintains efficient and robust support for user queries, providing precise and detailed explanations tailored to the educational datasets.

SUB-QUESTION 2: the Reynolds number: Re = ρVD/μ (use ρ = SG×1000, μ from cP→Pa·s, D = inside diameter in meters, V = Q/A with Q in m^3/s)
How would you solve this? Choose the best METHOD option (A/B/C/D):
A) Use the outside diameter directly (24 in) as D and plug into Re = ρVD/μ without converting viscosity.
B) Convert viscosity from cP to Pa·s, compute density from SG, compute inside diameter, compute Q in m^3/s, compute V=Q/A, then compute Re=ρVD/μ.
C) Compute Reynolds number as Re = V/D because density cancels for steady flow.
D) Use gallons/min and inches directly since Reynolds number is dimensionless.

Your choice (A/B/C/D): D
Why do you think this method is correct? reynolds number has no units so using gal/min and inches should still work.

Tutor feedback:
Not quite. Reynolds number is dimensionless, but you still must use consistent units inside the formula so they cancel correctly.
Use SI (or any consistent system) for ρ, V, D, and μ.
Correct Option: B

Now solve it.
Your final answer (number; you may include units): Re = 148,388
Your work/explanation (optional): Q=8750 gal/min *0.00378541 /60 = 0.55204 m^3/s; D_in=24-2*0.5=23 in=0.5842 m; A=pi*(D/2)^2; V=Q/A; rho=0.74*1000=740; mu=0.6 cP -> 0.006 Pa*s; Re=rho*V*D/mu = 148,388

Tutor feedback:
You're very close, but one conversion is off.
Check the cP → Pa·s step: you used a value 10× too large for μ, which makes Re 10× too small.

Try again.
Your final answer (number; you may include units): Re = 1.48388e6
Your work/explanation (optional): Same steps as above, but mu = 0.6 cP = 0.0006 Pa*s, so Re increases by 10x.

Tutor feedback:
Correct! Good catch—fixing viscosity scaling fixes the Reynolds number scaling.

Summary:
8/10 – method mistake first, then corrected your unit conversion and got the right order of magnitude.

==
☑ ALL SUB-QUESTIONS COMPLETED.

Overall summary:
Nice work overall. Your main slip was treating "dimensionless" as "any units are fine"—the formula still requires consistent units so they cancel correctly. Also watch small unit conversions (like cP to Pa·s), because a factor-of-10 mistake changes Reynolds number by the same factor. Final grade: 88/100.

Figure 7: Sub-question 2, providing support and guidance to the student when they have an error

Conclusion and Future Work

Our chatbot builds upon prior work from Reservoir Rock Properties and Fluid Flow and brings in interactive multiple-choice questions to help students solve longer engineering problems. This chatbot can serve as a resource for students who are struggling and considering dropping out of their majors. The bot is especially useful in courses with large section sizes when accessibility to teaching assistants is limited or not available.

Our future research work in engineering education will be to utilize the bot for further supporting student development in understanding the fundamentals related to this course. We also will investigate adding more sophisticated personalization in order that the bot can remember the errors that the students made and help support them so that they do not make the same mistake in follow-on problems. Longer term, we will test these and through IRB get approval to solicit student feedback and further improve the chatbot.

Artificial Intelligence Declaration Statement

This paper was written by authors at multiple institutions. When preparing the introduction and literature review section of the manuscript, the corresponding author utilized a secure version of Microsoft Copilot to help outline, organize and summarize key findings, as well as help write the first draft of this section in a consistent voice. After using this tool, the authors reviewed, edited and expanded this section. Furthermore, they reviewed and revised multiple versions of the entire work prior to finalizing it. The authors verified the accuracy of all statements and validity of materials referenced.

References

[1] Peyton, K., Unnkrishnan, S., and Mulligan, B. "A review of university chatbots for student support: FAQs and beyond," Discover Education, Vol. 4(21), 2025.

[2] Cunningham-Nelson, S. et al. "A review of chatbots in education: practical steps forward," 30th Annual Conference for the Australasian Association for Engineering Education (AAEE 2019): Educators Becoming Agents of Change: Innovate, Integrate, Motivate. 2019: Engineers Australia.

[3] El Mourabit, I., Jai Andaloussi, S., Ouchetto, O., and Miyara, M., "AI Chatbots in Higher Education: Opportunities and Challenges for Personalized and Mobile Learning," International Journal of Interactive Mobile Technologies, Vol. 19(12), 2025, https://doi.org/10.3991/ijim.v19i12.54163.

[4] Lee, L.-K., Fung, Y.-C., Pun, Y.-W., Wong, K.-K., Yu, M. T.-Y., and Wu, N.-I., "Using a multi-platform chatbot as an online tutor in a university course," in 2020 International Symposium on Educational Technology (ISET), 2020, pp. 53–56, https://doi.org/10.1109/ISET49818.2020.00021.

[5] Maderer, J. "Artificial Intelligence Course Creates AI Teaching Assistant," https://news.gatech.edu/news/2016/05/09/artificial-intelligence-course-creates-ai-teaching-assistant, May 2016, accessed January 2026.

[6] Chopra, S., Gianforte, R., and Sholar, J. "Meet Percy: The CS 221 Teaching Assistant Chatbot," ACM Transactions on Graphics, Vol. 1 (1), December 2016.

[7] Lluna, A. P. "Creation and Development of an AI Teaching Assistant," Master's Thesis, Universitat Politecnica de Catalunya, 2017/2018.

[8] Hubert. "6 Ways Artificial Intelligence and Chatbots Are Changing Education," Chatbots Magazine, May 2017.

[9] Gonda, D. E. and Chu, B. "Chatbot as a learning resource? Creating conversational bots as a supplement for teaching assistant training course," 2019 IEEE International Conference in Engineering, Technology and Education.

[10] Khan Academy, "Khanmigo for Learners," https://www.khanmigo.ai/learners, accessed January 2026.

[11] Liu, D. "How Sydney educators are building 'AI doubles' of themselves to help their students," https://educational-innovation.sydney.edu.au/teaching@sydney/how-sydney-educators-are-building-ai-doubles-of-themselves-to-help-their-students, December 2023, accessed January 2026.

[12] Hamam, D. "The New Teacher Assistant: A Review of Chatbots Use in Higher Education," in: Stephanidis, C., Antona, M., Ntoa, S. (eds) HCI International 2021 - Posters. HCII 2021. Communications in Computer and Information Science, vol 1421. Springer, Cham. https://doi.org/10.1007/978-3-030-78645-8_8.

[13] Sandu, N.,Gide, E. “Adoption of AI-chatbots to enhance student learning experience in higher education in India,” In: 2019 18th International Conference on Information Technology Based Higher Education and Training (ITHET), pp. 1–5. IEEE, September 2019.

[14] Haristiani, N., Danuwijaya, A. A., Rifai, M. M., and Sarila, H. “Gengobot: A chatbot-based grammar application on mobile instant messaging as language learning medium,” J. Eng. Sci. Technol. 14(6), 3158–3173 (2019).

[15] Deng, X., and Yu, Z. (2023). “A meta-analysis and systematic review of the effect of chatbot technology use in sustainable education,” Sustainability, 15(4), 2940.

[16] Ma, S. -P., Liang, Y. -C., Wang, S. -K., Huang, Y. -W and W. -L. You, "TABot: A Teaching Assistant Chatbot for Software Engineering Courses," 2023 30th Asia-Pacific Software Engineering Conference (APSEC), Seoul, Korea, Republic of, 2023, pp. 627-631.

[17] Schroeder, N. L. and Adesope, O. O. "A case for the use of pedagogical agents in online learning environments," Journal of Teaching and Learning with Technology, no. 2, pp. 43-47, 2013.

[18] Alsafari, B., Atwell, E. Walker, A., and Callaghan, M. “Towards effective teaching assistants: From intent-based chatbots to LLM-powered teaching assistants,” Natural Language Processing Journal, Volume 8, 2024,

[19] Arizona State University. Introducing sunny. [Online]. Available: https://heysunny.asu.edu/about, accessed January 2026.

[20] Georgia State University. Reduction of summer melt. [Online]. Available: https://success.gsu.edu/initiatives/reduction-of-summer-melt/, accessed January 2026.

[21] Sheffield University. Beacon - your digital guide. [Online]. Available: https://www.staffs.ac.uk/students/digital-services/beacon, accessed January 2026.

[22] Keith, J. M., Amirlatifi, A., Rahimi, S., Neupane, S., and Mittal, S. (2024, June), “Bark Plug: The ChatGPT of the Bagley College of Engineering at Mississippi State University,” Paper presented at 2024 ASEE Annual Conference & Exposition, Portland, Oregon. 10.18260/1-2—46635.

[23] Neupane, S., Hossain, E., Keith, J., Tripathi, H., Ghiasi, F., Goliarz, N. A., Amirlatifi, A., Mittal, S., and Rahimi, S. “From Questions to Insightful Answers: Building an Informed Chatbot for University Resources,” Paper presented at 2024 Frontiers of Education, Washington D.C.

[24] Keith, J. M., Amirlatifi, A., Rahimi, S., Mittal, S., Neupane, S., and Tripathi, H. (2025, June), “Barkplug 2.0 and Beyond: a Chatbot for Assisting Students in High DFW Courses,” Paper presented at 2025 ASEE Annual Conference & Exposition, Montreal, Quebec, Canada. 10.18260/1-2—55496.

[25] Loepp, E. “The Benefits of Higher-Order Multiple Choice Tests,” https://www.insidehighered.com/advice/2021/06/23/rethinking-multiple-choice-tests-better-learning-assessment-opinion, accessed January 2026.

[26] University of Michigan. “A Refresh on Writing Quality Multiple Choice Questions,” https://lsa.umich.edu/technology-services/news-events/all-news/teaching-tip-of-the-week/a-refresh-on-writing-quality-multiple-choice-questions.html, accessed January 2026.

[27] Bowen, R. “Understanding by Design,” https://cft.vanderbilt.edu/wp-content/uploads/sites/59/Understanding-by-Design.docx, accessed January 2026.

[28] Sotola, L.K. and Crede, M. “Regarding Class Quizzes: a Meta-analytic Synthesis of Studies on the Relationship Between Frequent Low-Stakes Testing and Class Performance,” *Educ Psychol Rev* **33**, 407–426 (2021). https://doi.org/10.1007/s10648-020-09563-9.

[29] Oc, Y., and Hassen, H. “Comparing The Effectiveness Of Multiple-Answer And Single-Answer Multiple-Choice Questions in Assessing Student Learning,” *Marketing Education Review*, *35*(1), 44–57 (2025). https://doi.org/10.1080/10528008.2024.2417106.

[30] Neupane, S., Mitra, S., Mittal, S., Golilarz, N. A., Rahimi, S., and Amirlatifi, A., “Medinsight: A multi-source context augmentation framework for generating patient-centric medical responses using large language models,” ACM Transactions on Computing for Healthcare, vol. 6(2), article 29(1-19) (2025), https://doi.org/10.1145/3709365.

[31] Mitra, S., Neupane, S., Chakraborty, T., Mittal, S., Piplai, A., Gaur, M., and Rahimi, S. “Localintel: Generating organizational threat intelligence from global and local cyber knowledge,” In: Adi, K., *et al.* Foundations and Practice of Security. FPS 2024. Lecture Notes in Computer Science, vol 15533. Springer, Cham. (2025) https://doi.org/10.1007/978-3-031-87496-3_5.