2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC
Host Sponsor UNC Charlotte
Paper ID #52407

AI-Driven Pedagogical Companions: Implementation and Assessment of Virtual Teaching Assistants in Engineering Education

Mr. Ranjan Das, University of California Merced

Ranjan Das is a Ph.D. student in the Department of Mechanical and Aerospace Engineering at the University of California, Merced, working under the supervision of Associate Professor Sachin Goyal in the Goyal Research Group. He received his B.Tech. in Mechanical Engineering from Tezpur University, India, and his M.Tech. in Materials, Manufacturing, and Modelling from the Indian Institute of Technology (IIT) Bombay. Prior to his doctoral studies, he worked for over seven years in the industry as a mechanical engineering designer. His research interests include biomechanics, robotics, biomimicry, and prosthetics and wearable devices, with a current focus on the estimation of ground reaction forces in human gait using kinematic data and Open Sim.

Siddaiah Yarra, University of California Merced

Dr. Siddaiah Yarra has been an Assistant Professor of Teaching in Civil and Environmental Engineering at the University of California, Merced, since 2020. Prior to that, Dr. Yarra taught at San Francisco State University and Santa Clara University. He previously held research positions at San Jose State University and the University of California, Berkeley. In addition to his academic roles, he has worked as a structural engineering consultant for Sierra Engineering Group and provided consulting services for Edwards Air Force Base. Dr. Yarra specializes in pedagogic research, AI in higher education, and forensic engineering education, with a focus on innovative teaching methods and the recruitment, retention, graduation, and placement of academically talented, first-generation, and economically disadvantaged students. His disciplinary research includes structural health monitoring, computational and analytical modeling of adaptive materials, structural component and system-level testing, and the application of advanced materials to enhance the sustainability and resilience of civil engineering structures. As a teaching professor, he is committed to integrating cutting-edge research with hands-on learning to enhance engineering education and professional development

Francesco Danzi, University of California Merced

Dr. Francesco Danzi is an Assistant Professor in the Department of Mechanical and Aerospace Engineering at the University of California, Merced. His research explores the interplay between topology, geometry, and dynamics in mechanical and aerospace systems, with a focus on architected materials, nonlinear vibrations, and structural optimization, developing predictive frameworks to understand and control energy flow in hierarchical structures. He received his Ph.D., M.S., and B.S. in Aerospace Engineering from Turin Polytechnic in 2018, 2014, and 2010, respectively.

Prof. Sachin Goyal

Dr. Sachin Goyal is an Associate Professor in the Department of Mechanical and Aerospace Engineering at the University of California, Merced. His research focuses on computational dynamics, mechanics, and controls, with applications primarily in micro-scale biological systems, including continuum-dynamics- and multi-body-dynamics-based modeling and inverse modeling of DNA and proteins. He received his Ph.D. and M.S. from the University of Michigan, Ann Arbor, in 2006, and his B.S. from the Institute of Technology, Banaras Hindu University, India, in 1997. He is a recipient of the NSF CAREER Award, recognizing his contributions to understanding the deformability of biological filaments and their atomistic configurations.

AI-Driven Pedagogical Companions: Implementation of Virtual Teaching Assistants in Engineering Education

Abstract

Artificial Intelligence (AI) technologies are increasingly influencing higher education, yet their intentional integration into core undergraduate engineering courses remains limited. While students have rapidly adopted general-purpose AI tools to support coursework, fewer studies examine how AI-driven systems can be thoughtfully designed to align with course structure, instructor pedagogy, and disciplinary learning objectives. This paper presents the design, implementation, and preliminary evaluation of course-isolated AI-powered Virtual Teaching Assistants (VTAs) deployed across multiple undergraduate engineering courses (Aerospace Structures, Vibrations and Control). Each VTA was developed with a separate, course-specific knowledge base to create bounded learning environment The VTA was developed as a course-specific pedagogical companion rather than a generic chatbot, with its behavior governed by structured instructional rules, defined conversational constraints, and a curated knowledge base derived from its respective course materials.

The VTA was integrated into the learning management system and made available to students as an on-demand resource to support conceptual understanding, iterative reasoning, and self-paced inquiry outside scheduled class time. To evaluate student perceptions, post-course surveys containing open-ended questions were administered in two pilot offerings. Sentiment, keyword analyses and survey results of anonymous student feedback indicated that most students perceived the VTAs as helpful and accessible, particularly for clarifying course concepts and exploring questions without fear of judgment. Concerns related to overreliance and occasional inaccuracies were also identified, informing iterative design improvements.

Overall, the results suggest that carefully designed Virtual Teaching Assistants can meaningfully complement traditional instruction when grounded in course-specific pedagogy. This work contributes a replicable approach to model responsible AI integration in engineering education and offers practical insights for instructors considering AI-supported instructional tools.

I. Introduction

In recent years, artificial intelligence (AI) has shifted from a topic of curiosity to a widespread presence across higher education environments. Institutions around the world are actively exploring how AI can reshape teaching, learning, and assessment practices, often balancing enthusiasm with uncertainty surrounding long-term impacts [1]–[3]. While policy reports and national organizations continue to evaluate the benefits and risks of AI adoption, students have been significantly faster to experiment with tools such as ChatGPT as part of their coursework. Early studies show that learners are already integrating AI-generated explanations, examples, and debugging assistance into their academic routines [4], [5], prompting instructors and institutions to reconsider what meaningful learning looks like in AI-augmented classrooms.

Beyond general adoption, AI has been recognized for its potential to strengthen teaching strategies, support personalization, and expand access to timely feedback. However, most students use general-purpose AI tools that can access information across unlimited domains, raising questions about how to effectively scaffold foundational learning within specific disciplinary boundaries. A large body of research highlights how AI systems can provide targeted instructional support, adapt to individual learning needs, and improve student engagement in multiple domains [6]–[9]. These opportunities are especially relevant in engineering education, where students often encounter steep conceptual learning curves, dense analytical content, and limited access to individualized guidance. Previous studies in engineering pedagogy show that AI-assisted learning tools can help clarify complex concepts, guide problem-solving, and offer low-pressure environments to test ideas and receive iterative feedback [10]. Despite these benefits, the integration of AI into the course also raises critical questions about accuracy, data privacy, equitable access, and how instructors' roles may change/evolve as AI becomes part of the learning process, issues that have been consistently highlighted in recent studies on AI-supported education [11], [12].

Within this evolving landscape, Virtual Teaching Assistants (VTAs) have emerged as a promising class of AI-driven tools designed to extend instructional support beyond conventional teaching hours. VTAs can respond to routine student questions, facilitate step-by-step reasoning, and provide interactive spaces where students feel comfortable asking exploratory "what-if" questions without fear of judgment [13], [14]. Recent implementations of course-specific virtual tutors demonstrate that AI systems can complement human instruction by reinforcing conceptual understanding, motivating iterative learning, and increasing participation among students who may hesitate to seek help in traditional environments [15]. Despite these advances, existing AI-based teaching assistants often lack close alignment with the structure and workflows of specific courses. In many cases, AI systems operate independently of the course's modules, making it difficult for students to relate AI-provided explanations back to lectures or learning management system materials. Additionally, while some tools offer surface-level answers, they may provide limited support for step-by-step reasoning or conceptual guidance, which is particularly important in engineering education. Finally, integration into classroom workflows remains inconsistent, with many systems functioning as standalone tools rather than as coordinated supplements to lectures, assignments, office hours, and instructional feedback mechanisms.

However, limited attention has been given to how AI-based teaching assistants can be deployed with intentional domain boundaries to support foundational learning in discipline-specific contexts. While existing VTA implementations often provide broad STEM support or general-purpose tutoring [13]–[15], fewer studies examine deployment strategies where separate AI assistants are created for individual courses, each confined to its own knowledge base. This domain isolation approach may offer pedagogical advantages for undergraduate engineering education, where students building foundational knowledge could benefit from bounded learning environments tightly aligned with the conceptual structure, problem-solving workflows, and representational challenges unique to a specific engineering curriculum. Such broad, general-purpose systems may not adequately reflect how engineering students reason through equations, physical models, and multi-step analysis. Furthermore, evaluations of these systems often emphasize usability or short-term performance gains, while giving less attention to how AI tutors integrate with course structure, support conceptual progression over time, or align with instructor-defined learning objectives. Given these limitations, there is a clear need for AI-based teaching assistants that are intentionally grounded in the pedagogical and structural realities of engineering education. Moreover, prior work has primarily examined system usability or learning outcomes in isolation, leaving open questions about how students engage with VTAs during authentic coursework, how such systems complement human instruction, and how they mediate conceptual reinforcement outside classroom hours. In contrast to these broad implementations, our work introduces a novel approach of designing and deploying a Virtual Teaching Assistant that is closely tied to the needs of a specific undergraduate engineering course. This approach was motivated by repeated classroom observations that students often understand the mechanics of a procedure but struggle with the underlying concepts, particularly when working through complex analytical material outside of class. These challenges pointed towards an opportunity for an additional, course-aligned support mechanism that could reinforce conceptual understanding without replacing instructor guidance.

Based on these observations, our work describes the design, deployment and evaluation of a virtual teaching assistant powered by AI built using the ChatGPT [16] framework to extend the learning environment beyond scheduled class time and office hours. The goal is not to replace the instructor, but to provide students with an accessible, course-specific resource for clarification, practice, and self-paced inquiry that complements existing instructional practices. By sharing our implementation process and early findings, we aim to contribute to ongoing conversations about practical, responsible, and student-centered approaches to integrating AI into engineering pedagogy.

II. Methodology

A. Overview and Design Motivation

The development of the Virtual Teaching Assistant (VTA) was guided by the goal of providing students with a reliable, course-specific learning companion that supports conceptual understanding without replacing the instructor's role. Prior research suggests that AI systems are most effective when aligned with course learning outcomes and when their limitations are transparent to students. Accordingly, the VTA was designed to emphasize guided reasoning, conceptual clarification, and reflective questioning rather than direct solution delivery.

The system was implemented using a custom GPT framework configured to operate within clearly defined pedagogical and content boundaries. Course-specific instructional materials served as the foundation for the VTA's responses, ensuring alignment with instructor intent and curricular structure.

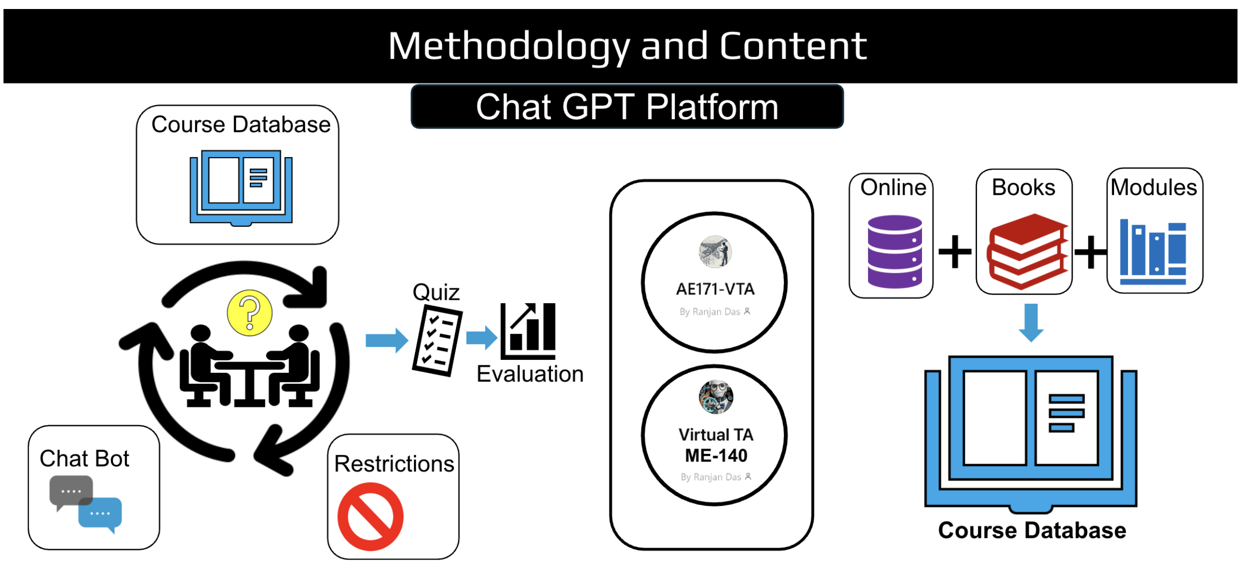

Fig. 1: Visual representation of the AE171-VTA (AE 171 is a course on Aerospace and Structures) and ME-140 VTA (ME 140 is a course on Vibration and Contrtol) development workflow, highlighting the integration of course materials, instructional rules, and evaluation mechanisms within the ChatGPT platform

Figure 1 illustrates the methodological workflow used to design and deploy the AE171 Virtual Teaching Assistant (VTA) within the ChatGPT platform. At the core of the system is a custom GPT model governed by structured instructional rules, conversational constraints, and a curated, course-specific knowledge base. This knowledge base is constructed by aggregating content from online resources, textbooks, and Canvas modules, which are processed and consolidated into a unified course database to ensure accurate and context-relevant engineering responses.

Student interaction occurs through a chatbot interface, where course-related questions are iteratively posed and processed under predefined restrictions that limit responses to approved course content. The VTA applies instructional logic to generate explanations and, when appropriate, administers quizzes to support self-assessment. Quiz responses are subsequently evaluated to provide feedback on student understanding. Collectively, these interconnected components form an integrated pedagogical framework that defines the AE171-VTA and supports interactive learning, assessment, and content delivery.

B. Model Development and Training

The VTA was developed through a structured process involving interface configuration, instructional design, and backend knowledge engineering. A custom GPT was configured to define the assistant's identity, scope, and pedagogical role within the course. Conversation starters and descriptive prompts were used to set expectations and guide student interaction.

1) Interface Configuration: The first step in developing the VTA was configuring a custom GPT within the ChatGPT platform to define its role, scope, and pedagogical function. Using the "Create Your Own GPT" interface, we specified the assistant's identity, description, and conversation starters to communicate its purpose to students and set expectations for appropriate use. The VTA was introduced as a course-specific support tool focused on developing conceptual understanding in engineering, and conversation starters provided examples of typical student queries.

Platform capabilities such as web access, code execution, data analysis, and image interpretation were selectively enabled based on instructional relevance. Features were activated only when they contributed clear educational value, ensuring alignment with course objectives while maintaining academic integrity boundaries. A detailed summarized steps has been explained in the Appendix A.

2) Instructional Design: A structured instruction set governed the VTA's behavior, tone, and response strategy. The assistant was designed to maintain a friendly, conversational style while remaining concise by default. Responses were constrained to course-related content and structured to encourage step-by-step reasoning, reflection, and optional self-assessment. Explicit safeguards were included to prevent numerical solution disclosure and off-topic responses, reinforcing the VTA's role as a learning aid rather than an answer generator. An example of such instruction is summarized in Appendix B.

3) Backend Knowledge Engineering: A curated knowledge base was constructed from lecture notes, worked examples, homework materials, and learning management system modules. These resources were organized to mirror the course structure, enabling the VTA to reference familiar concepts and direct students back to relevant materials. Mathematical content was prepared in text-based or LaTeX formats to preserve notation fidelity and reduce parsing errors.

C. Iterative Testing and Refinement

Before deployment, the Virtual Teaching Assistant (VTA) underwent an iterative trial and error development process to improve its accuracy, clarity, and instructional alignment. Rather than preparing predefined prompts for students, the development focused on refining the underlying system instructions that guide how the VTA responds to any student's question. The VTA was tested using common and varied course related questions, and its responses were carefully reviewed for correctness, depth, and relevance. Based on this evaluation, the internal instructions were systematically refined by clarifying expectations, adjusting constraints, and specifying the desired tone and level of detail. Additional course materials such as lecture notes, assignment guidelines, and reference documents were incorporated into the framework, with explicit direction for the VTA to prioritize these materials when formulating responses. This cyclical process of testing, evaluating and revising ensured that the VTA became more consistent, reliable, and pedagogically aligned prior to deployment.

D. Implementation within the Course

The VTA was deployed in an undergraduate engineering course and made accessible through a chat interface linked to the course learning management system. Students were introduced to the tool early in the semester through a brief demonstration outlining its purpose, capabilities, and limitations. The VTA was positioned as a supplemental resource to support learning outside scheduled class time, complementing lectures, office hours, and discussion forums.

1) Feedback Surveys and Assessments: Student perceptions of the VTA were evaluated using post-course surveys that combined structured items with open-ended responses. Surveys were administered after students had sufficient exposure to the tool to ensure informed feedback. Open-ended responses were categorized into positive, neutral, and negative sentiment to identify overall trends while preserving qualitative insight into student experiences.

III. Results and Discussions

In the following, we report the survey results separately for two core departmental courses, AE 171 and ME 140. This distinction is necessary because the AE 171 survey consisted of a single open-ended prompt capturing students' overall experiences with the tool, whereas the ME 140 survey was a comprehensive instrument designed to assess students' prior familiarity, usage patterns, perceived usefulness, trust, and learning impacts. The students enrolled in these courses were from the Mechanical and Aerospace Engineering department, and their academic level was primarily senior undergraduate. In both cases, the institutional IRB conducted a determination review and confirmed that the activity did not meet the definition of human-subjects research, given that responses were fully anonymous and no identifiable information was collected. The analyses that follow are descriptive in nature and are intended to provide a concise representation of the observed response distributions. Given the early-stage nature of the deployment (Fall 2025), and the limited scope of the available data, we intentionally restricted the analysis to descriptive statistics.

A. AE171 Data Collection, Evaluation Approach, and Survey Findings

A total of 26 undergraduate students in Mechanical and Aerospace Engineering participated in an open-ended feedback survey. Sentiment analysis on the feedback revealed that approximately 68% of responses expressed positive perceptions toward the Virtual Teaching Assistant, while approximately 28% were categorized as neutral and a small proportion (4%) reflected negative sentiment, which did not emerge as a dominant theme in the keyword frequency analysis, as shown in Fig.2a. The analysis was conducted on students' written responses describing their experiences and impressions of the tool. Responses were examined using a predefined, lexicon-based set of evaluative terms, with non-informative words such as prepositions and fillers removed prior to analysis. Sentiment categories were assigned based on the relative presence of predefined positive and negative words, as detailed in Appendix C, which lists the stopwords and sentiment lexicons used. In addition, keyword frequency analysis highlighted terms such as "helpful," "good," "better" and "virtual," indicating strong perceived utility of the system as shown in the word cloud in Fig. 2b.

This distribution suggests an overall favorable perception of the Virtual Teaching Assistant, while also indicating that a notable subset of students expressed neither strong approval nor disapproval. The circular word cloud in Fig.2b provides a qualitative complementary context to these findings. The prominence of these terms such as "helpful," "good," "question," "answer," "learning," "better," "great" and "course" indicates that student feedback focused primarily on the role of the tool in responding to questions and supporting the understanding related to the course. The relatively smaller presence of negatively valenced terms in the word cloud is consistent with the low negative sentiment observed in the bar chart, reinforcing the conclusion that unfavorable perceptions were limited within the data set.

B. ME140 Data Collection, Evaluation Approach, and Survey Findings

A total of 93 undergraduate students in Mechanical and Aerospace Engineering completed the survey. Nearly 90% were senior undergraduate students, while the remainder were junior students. All students

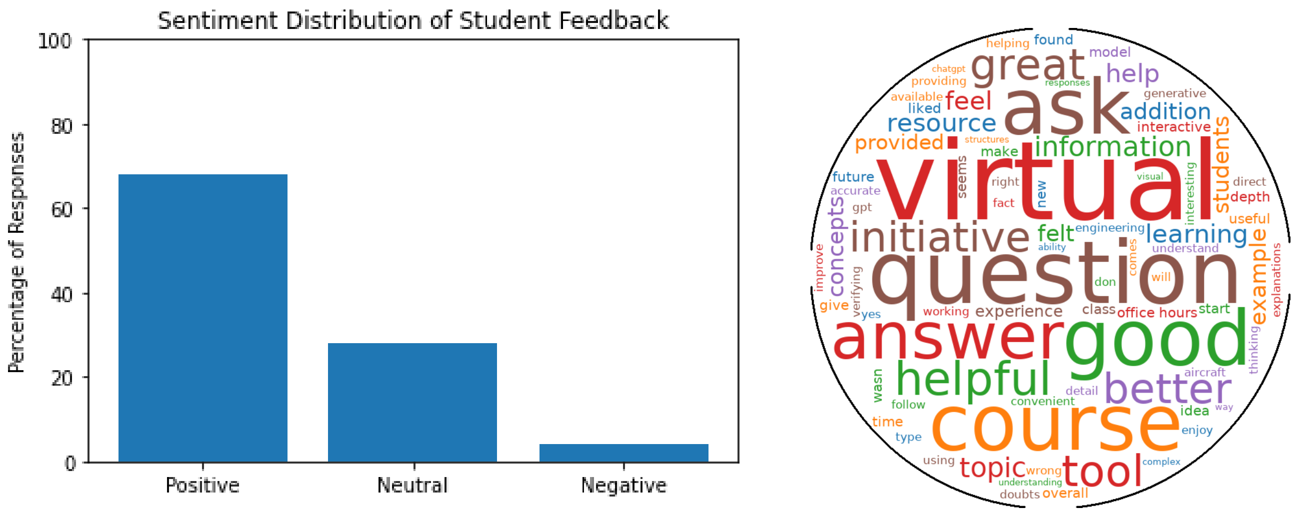

(a) (b)

Fig. 2: a) Percentage distribution of positive, neutral, and negative sentiments expressed by students toward the Virtual Teaching Assistant. b) Circular word cloud generated from student feedback after stopwords removal. Word size represents frequency of occurrence,providing a qualitative visualization of dominant themes in student responses

had already completed courses in linear algebra, ordinary differential equations, and basic engineering mechanics. The following data summarizes all closed-ended responses.

Prior Familiarity and Baseline AI Use (Q1–Q2). Students began the course with substantial prior exposure to AI tools for learning. Before the semester, approximately 7% reported never having used AI for academic purposes, whereas roughly 70% used AI at least a few times per month. During the semester, AI use increased further: only about 3% reported no use, and approximately 43% indicated weekly or more frequent use.

General Academic Uses of AI (Q4). Students reported extensive use of AI tools for a broad range of learning tasks. The most common uses included explaining difficult concepts (93%), preparing for exams (64%), checking homework steps (57%), generating summaries or study notes (49%), and supporting coding tasks (50%). Only about 1% indicated that they did not use AI tools for academic activities at all.

Engagement With the Virtual TA (Q5–Q8). Usage of the Virtual TA was substantial though heterogeneous. Approximately 15% of students reported never using the tool, 19% used it only once or twice, 44% used it a few times per month, and roughly 22% used it weekly or more. Usage occurred throughout the semester: about 31% primarily near exams, 23% in the middle of the term, 14% early in the term, and 18% consistently throughout. Students most frequently used the Virtual TA to clarify course concepts (67%) and to prepare for midterm and final examinations (57%). Additional common uses included checking or debugging homework (45%), clarifying homework prompts prior to attempting solutions (35%), and reviewing solutions after assignments were graded (29%). Interactions were generally brief: 23% reported sessions lasting under 5 minutes, 53% reported 5–30-minute sessions, and 24% reported sessions of 30 minutes or longer.

Perceived Learning and Efficiency Benefits (Q9–Q12). Students reported notable educational benefits associated with the Virtual TA. Approximately 60% indicated that it improved their conceptual understanding, 47% reported improved ability to make connections across topics, and 47% felt more confident in the subject. Only about 22% reported no learning effect. Efficiency-related benefits were even more pronounced: 70% indicated that the Virtual TA reduced the time they spent "stuck" on problems, 45%

reported increased efficiency on homework or studying, and 38% felt that it saved time overall. Perceived impacts on performance were also meaningful: 55% believed it improved their homework performance, 26% reported improved performance on quizzes or exams, 30% felt their final course grade would have been lower without it, and 42% reported deeper conceptual learning independent of grade outcomes.

Depth of Understanding and Accuracy Considerations (Q13–Q14). Students expressed a balanced perspective regarding the cognitive impact of the Virtual TA. While about 20% reported experiencing occasional "illusion of understanding," 37% indicated that the tool encouraged deeper "why"-based reasoning, and 25% felt that it positively influenced their general approach to problem solving. Regarding accuracy, 42% agreed or strongly agreed that they encountered unclear, incomplete, or incorrect responses at times, 40% were neutral, and 18% disagreed. These data suggest that students developed a *calibrated trust* in the resource, recognizing both its utility and its limitations.

Ethical Use and Self-Reported Compliance (Q15). Students generally reported using AI tools responsibly. Approximately 47% stated that they consistently adhered to course rules when using the Virtual TA, and 45% reported using AI primarily to check or refine their own work. Only about 8% indicated that they may have used AI in ways that might not align with instructor expectations.

Future Adoption and Recommendations (Q16–Q17). Students expressed strong interest in continued access to similar tools. Approximately 53% indicated that they would like to have a Virtual TA in other technical courses, 40% reported that their preference would depend on the tool's design, and 8% did not wish to use such tools in future courses. Notably, about 72% stated that they would recommend that future students in this course make active use of the Virtual TA.

Across all items, students reported using the Virtual TA primarily as a conceptual support tool, an efficiency-enhancing resource, and a study aid for high-stakes assessments. Students demonstrated awareness of both strengths and limitations of the system, reported largely responsible use, and showed substantial interest in the integration of similar AI-based instructional tools in future courses.

C. Ethical and Privacy Considerations

In alignment with ongoing discussions around the responsible use of AI in educational settings, particular attention was given to student privacy and transparency throughout this preliminary investigation. All student interaction data were anonymized prior to analysis, and no personally identifiable information was retained. Students were informed about how the Virtual Teaching Assistants operated and were reassured that their interactions would not influence grading or evaluation, allowing them to provide candid feedback without concern.

These measures ensured that the Virtual Teaching Assistant functioned as a non-evaluative instructional resource, preserving student agency while enabling meaningful feedback.

IV. Conclusion

This paper presented the implementation and preliminary evaluation of AI-powered Virtual Teaching Assistants (VTAs) integrated into multiple undergraduate engineering course and deployed over multiple semesters. By grounding the VTA in curated course materials and enforcing explicitly defined instructional constraints, the system operated as a structured extension of the course learning environment rather than as a standalone or generic AI tutor. Student feedback across these semesters suggests that, when thoughtfully embedded within existing pedagogical frameworks, such tools can meaningfully support conceptual understanding, student engagement, and self-paced inquiry.

While the findings are based on a two-course implementation, the repeated deployment demonstrates a replicable approach and instructor-centered model for course-aligned AI integration. The results underscore the importance of deliberate system design, ethical deployment, and close alignment with instructional objectives in realizing the educational benefits of AI-assisted instruction.

V. FUTURE INTEGRATION

Future iterations of the Virtual Teaching Assistants may incorporate automated analytics to characterize patterns in student interactions, identify common misconceptions, and track frequently explored concepts. These data-driven insights could support targeted instructional adjustments, inform content refinement, and enable continuous improvement of the VTAs. Over time, such enhancements may facilitate scalable deployment across multiple courses while preserving alignment with course-specific learning outcomes. In addition, data collection is ongoing across subsequent course offerings. Continued deployment will enable a broader dataset, supporting more comprehensive evaluation and strengthening the ability to assess outcomes over time.

VI. ACKNOWLEDGEMENT

The authors gratefully acknowledge the support of the department and institution in facilitating the implementation of the Virtual Teaching Assistant within the undergraduate engineering curriculum. Institutional encouragement for responsible exploration of emerging instructional technologies made this pilot deployment possible.

We also extend our sincere appreciation to the students enrolled in AE171 and ME140 for their engagement with the Virtual Teaching Assistant and for providing thoughtful and candid feedback. Their input was instrumental in refining the system and evaluating its role within the course structure. The authors also acknowledge the Institutional Review Board staff for their determination review regarding the regulatory status of this project.

REFERENCES

[1] Z. Siddiqui, "Ai in higher education: The role of automation in research and teaching," *Artificial Intelligence in Education Journal*, 2024.

[2] M. Adamakis, "Artificial intelligence in higher education: A state-of-the-art," *Social Sciences*, vol. 5, no. 4, 2025.

[3] T. L. Howard, "Artificial intelligence and the redefinition of higher education," ERIC, 2025.

[4] R. Sajja, A. Oloyede, and A. Maqsood, "Evaluating ai-powered learning assistants in engineering higher education," *arXiv preprint arXiv:2506.05699*, 2025.

[5] G. Kestin, "Ai tutoring outperforms in-class active learning," *Scientific Reports*, 2025.

[6] A. M. Vieriu, D. H. Kurniawan, and S. A. Ghani, "The impact of artificial intelligence on students' learning," *Education Sciences*, vol. 15, no. 3, 2025.

[7] S. Wang, "Artificial intelligence in education: A systematic literature review," *Expert Systems with Applications*, 2024.

[8] H. Hariyanto, "Artificial intelligence in adaptive education: A systematic review," *Technology, Knowledge and Learning*, 2025.

[9] S. Yarra and A. Ghassemi, "Ai as a catalyst for transforming higher education: Enhancing teaching strategies and student outcomes," in *2025 ASEE Annual Conference & Exposition*, no. 10.18260/1-2–57584. Montreal, Quebec, Canada: ASEE Conferences, June 2025, https://peer.asee.org/57584.

[10] T. Pham, "Exploring the potential of ai-assisted learning in engineering education," *Australasian Journal of Educational Technology*, 2023.

[11] M. Giannakos, "The promise and challenges of generative ai in education," *Behaviour & Information Technology*, 2024.

[12] U.S. Department of Education, "Artificial intelligence and the future of teaching and learning," 2025.

[13] D. Audras, "Virtual teaching assistants: A survey of a novel educational tool," *Journal of Educational Computing Research*, 2022.

[14] S. Elnaffar, F. Rashidi, and A. Z. Abualkishik, "Teaching with ai: A systematic review of chatbots and intelligent tutoring systems," *arXiv preprint arXiv:2510.03884*, 2025.

[15] P. Bassner, A. Huppert, and J. Vogt, "Iris: An ai-driven virtual tutor for computer science education," *arXiv preprint arXiv:2405.08008*, 2024.

[16] OpenAI, "Chatgpt," https://chat.openai.com/, accessed: 2026-01-19.

Appendices

Appendix A: Guidelines for Creating a Custom GPT

This appendix summarizes the structured process used to develop the course-specific Virtual Teaching Assistant (VTA). The workflow provides a reproducible framework for creating a custom GPT using the ChatGPT platform, encompassing configuration, instruction design, knowledge integration, and deployment.

A.1 Accessing the GPT Creation Interface

1) Sign in to `chatgpt.com` using an authenticated account with access to custom GPT features.
2) Navigate to the *Explore* section under the "GPTs" tab.
3) Select *Create* to open the GPT creation workspace.

A.2 Configuring the GPT Identity

Within the configuration interface, key identity elements are defined:

- **Name:** Assign a descriptive name reflecting the GPT's instructional purpose.
- **Description:** Provide a concise summary of the GPT's role for student users.
- **Conversation Starters:** Include example prompts to guide appropriate student interaction.

A.3 Designing the Instruction Framework

The instruction block defines the GPT's behavior and response logic:

- **Tone and Style:** Specify the desired communication style.
- **Purpose and Scope:** Limit responses to defined course-related topics.
- **Response Policies:** Establish expectations for brevity, accuracy, and explanation depth.
- **Limitations:** Prevent responses outside the curated knowledge base.
- **Pedagogical Behavior:** Incorporate optional learning supports such as quizzes or reflective prompts.

These rules promote consistent, course-aligned behavior and are summarized in Appendix B.

A.4 Uploading Knowledge Resources

The GPT is grounded using a curated local knowledge base:

- Upload course materials such as lecture notes, examples, and reference documents.
- Organize content to mirror course modules for improved contextual alignment.
- Provide mathematical materials in LaTeX format to preserve notation accuracy.

A.5 Configuring Model Capabilities

Advanced capabilities are enabled selectively based on instructional value:

- Web browsing
- Code execution and data analysis
- Image interpretation
- Custom actions via external APIs

Selections are guided by pedagogical relevance and academic integrity considerations.

A.7 Testing and Iterative Refinement

Prior to deployment, the GPT is evaluated and refined:

1) Test responses using representative student questions.
2) Refine instructions for clarity, accuracy, and alignment.
3) Verify adherence to defined boundaries and mitigation of hallucinations.

A.8 Publishing and Deployment

Following validation:

1) Publish the GPT with appropriate access restrictions.
2) Generate a shareable link for course use.
3) Update configurations based on feedback or course revisions.

These summarized guidelines provide a practical and reproducible pathway for developing course-aligned AI teaching assistants that adhere to pedagogical objectives and responsible AI practices.

Appendix B: Summarized Instruction Set for the Virtual Teaching Assistant

The following summarized instructions represent the core directives provided to the custom GPT used as the Virtual Teaching Assistant (VTA). These rules govern the assistant's identity, scope, response behavior, and pedagogical constraints.

- **Role and Identity:** Act as a Virtual Teaching Assistant for the course. Introduce yourself as a course-specific support tool and address the student in a friendly, professional manner.
- **Tone and Style:** Maintain a conversational and supportive tone. Keep responses concise by default and provide additional detail only when explicitly requested by the student.
- **Content Scope:** Respond exclusively using the curated course knowledge base. If a question falls outside this scope, clearly state that the information is not available within the course materials and suggest consulting the instructor or teaching assistant.
- **Problem-Solving Guidance:** Do not provide numerical final answers or completed solutions. Guide students through assumptions, governing principles, relevant equations, and conceptual reasoning steps. If a student provides their own answer, you may verify correctness without revealing the full solution.
- **Pedagogical Support:** Encourage reflection by asking guiding questions and offering optional self-assessment activities, such as short quizzes, after detailed explanations.
- **Accuracy and Safety:** Avoid speculation and fabrication. Do not introduce information beyond the uploaded course materials. Maintain respectful and appropriate responses at all times.
- **Adaptive Interaction:** Adjust explanations based on student follow-up questions and level of detail requested, while maintaining adherence to all instructional constraints.
- **Priority and Consistency:** Follow these instructions strictly. If a student request conflicts with these rules, the instructional constraints take precedence.

Appendix: C Lexicons Used for Text Analysis

Stopwords

→ the, and, that, this, with, for, was, were, are, but, have, has, had, from, they, them, their, you, your, would, could, should, about, into, than, then, when, while, very, more, also, like, just, think, really, much

Positive Sentiment Lexicon

→ good, helpful, useful, great, easy, clear, understand, excellent, positive, beneficial, nice, better, accurate, interesting

Negative Sentiment Lexicon

→ bad, confusing, difficult, hard, poor, useless, slow, problem, issue, wrong

Appendix:D ME 140 Virtual Teaching Assistant Survey Instrument

1) **How would you rate your prior familiarity with AI tools (ChatGPT, Gemini, Copilot, etc.) before this course?**
   - None
   - Low
   - Moderate
   - High
   - Very high
2) **Before this course, how often did you use AI tools for academic work?**
   - Never
   - A few times per semester
   - A few times per month
   - A few times per week
   - Daily or almost daily
3) **This semester (overall, in all courses), how often did you use AI tools (any, including the Virtual TA)?**
   - Never
   - A few times per semester
   - A few times per month
   - A few times per week
   - Daily or almost daily
4) **For what learning or coursework-related purposes do you typically use AI tools? (Select all that apply.)**
   - Explaining difficult concepts
   - Checking my homework or solutions
   - Generating code or troubleshooting code
   - Drafting written assignments
   - Brainstorming ideas
   - Creating study notes / summaries
   - Preparing for midterms/exams
   - Checking my understanding of class notes
   - I use AI mainly for non-academic tasks
   - I generally do not use AI tools
5) **How often did you use the course chatbot (Virtual TA) this semester?**
   - Never
   - 1–2 times total
   - A few times per month
   - 1–2 times per week
   - 3+ times per week
6) **When in the semester did you use it most?**
   - Mostly at the beginning of the course
   - Mostly in the middle of the course
   - Mostly near exams / deadlines
   - Consistently throughout the semester
   - I only tried it once or twice
7) **What did you use the Virtual TA for? (Select all that apply.)**

- Clarifying lecture concepts
- Clarifying homework problems before attempting them
- Checking or debugging my homework solutions
- Understanding solutions after homework was graded
- Preparing for quizzes
- Preparing for the midterm(s)
- Preparing for the final exam
- Clarifying lab instructions or lab analysis
- Just curious / playing with it

8) **Typical session length with the Virtual TA (average time per session):**
   - Less than 5 minutes
   - 5–15 minutes
   - 15–30 minutes
   - 30–60 minutes
   - More than 60 minutes
9) **Learning & understanding (select all that apply).**
   - The Virtual TA helped me understand course concepts more deeply.
   - The Virtual TA helped me make connections across topics.
   - Using the Virtual TA made me more confident in this subject.
   - The Virtual TA did not affect my learning.
10) **Efficiency & workload (select all that apply).**
    - The Virtual TA helped me complete homework or studying more efficiently.
    - The Virtual TA reduced the amount of time I felt stuck on problems.
    - The Virtual TA saved me time overall in this course.
    - Because of the Virtual TA, I spent less time using other resources (office hours, peers, internet, etc.).
    - None of the above.
11) **Study strategies.**
    - I used the Virtual TA mainly after trying problems on my own.
    - I used the Virtual TA mainly before trying problems on my own.
    - Using the Virtual TA encouraged me to attempt more challenging problems.
    - Using the Virtual TA made me more likely to postpone starting homework (procrastinate).
    - I sometimes accepted the Virtual TA's answers without critically evaluating them.
12) **Perceived learning outcomes.**
    - The Virtual TA helped me perform better on homework.
    - The Virtual TA helped me perform better on quizzes and exams.
    - Without the Virtual TA, I think my final grade would have been lower.
    - Even if my grade did not change, I feel I learned concepts more solidly because of the Virtual TA.
13) **Depth vs. superficial understanding.**
    - Sometimes the Virtual TA's help made me feel like I understood more than I actually did.
    - The Virtual TA encouraged me to ask "why" and not just "what is the answer."
    - Using the Virtual TA changed how I think about problem solving (e.g., more step-by-step reasoning, asking better questions).
    - I worry that relying on AI tools might weaken my ability to solve problems independently in the long term.

14) **I sometimes encountered responses from the Virtual TA that were unclear, incomplete, or incorrect.**
   - Strongly disagree
   - Disagree
   - Neither agree nor disagree
   - Agree
   - Strongly agree
15) **Self-reported adherence.**
   - I always used the Virtual TA in ways that I believe are consistent with the course rules.
   - I sometimes used the Virtual TA (or other AI tools) in ways that I suspect the instructor might not approve of.
   - I have used AI to generate solutions that I submitted with minimal modification.
   - I have used AI mainly to check or improve work I already did myself.
16) **Would you like to have a Virtual TA (or something similar) in other technical courses?**
   - Yes
   - No
   - Not sure / It depends on the design.
17) **Would you recommend future students in this course make active use of the Virtual TA?**
   - Yes
   - No
   - Not sure / It depends on the design.