2026 ASEE ANNUAL CONFERENCE June 21–24, 2026 | Charlotte Convention Center, NC

Host Sponsor UNC Charlotte

Paper ID #53358

Exploring the Impact of AI-Based Assistants on Student Learning in Engineering Courses: A Thermodynamics Case Study (Work in Progress)

Dr. Zahra Pournorouz, Stevens Institute of Technology (School of Engineering and Science)

”Zahra Pournorouz received her Bachelor of Science degree in Aerospace Engineering from Amirkabir University of Technology (Tehran Polytechnic) in Tehran in the Fall of 2014. After finishing her bachelor’s studies, she got admitted directly to the Ph.D. program in Mechanical Engineering at the University of Texas at Arlington and graduated in August 2018. Her research interests mainly focus on oil-based nanofluids and enhancing the thermophysical properties of synthetic oils. This was the first demonstration of the work ever done in this field and resulted in broad environmental and cost benefits, especially in energy storage and heat transfer applications. She has more than three years of experience teaching thermofluidic, mechanical design, and solid and structure courses and supervising senior capstone projects collaborating with industries such as Saint-Gobain, Klein Tools, and Parker. She also has served in leadership roles at the Society of Women Engineers and STEM advisory task force to represent diversity and inclusion and improve student success and retention for underrepresented students.”

Exploring the Impact of AI-Based Assistants on Student Learning in Engineering Courses - A Thermodynamics Case Study

Abstract

This Work-in-Progress paper reports the design and implementation-planning phase of an AI-based teaching assistant intended to support student learning in Engineering Thermodynamics. Thermodynamics remains challenging for many undergraduate engineering students because it requires students to coordinate abstract concepts, property data, mathematical models, and system-level reasoning. The project was originally intended to include a classroom deployment during the current course offering; however, the AI tutor was not deployed to students because development is still underway and Institutional Review Board (IRB) approval for student testing and data collection is still pending. Therefore, no students used the tool as part of a study, and this paper does not report student outcome data, usage data, or student perceptions. Instead, it documents the instructional motivation, design decisions, prototype architecture, responsible-use planning, academic-integrity considerations, and assessment framework that will guide a future classroom study after approval is obtained. The proposed AI assistant is designed to provide scaffolded explanations, guided problem-solving support, and multiple representations of thermodynamic concepts while preserving the role of human instruction. The planned future study, to be implemented only after IRB approval and tool validation are complete, will compare a section using the AI assistant with a section receiving standard instruction, using concept-inventory data, course performance measures, usage analytics, surveys, and qualitative feedback. By reframing the delayed deployment as a design and implementation finding, this WIP paper contributes practical guidance for engineering educators seeking to develop generative-AI tools responsibly before introducing them into required technical courses.

Introduction

Engineering Thermodynamics is a core course in many undergraduate engineering curricula. At Stevens Institute of Technology, the course serves students from multiple engineering and applied-science majors, many of whom do not identify primarily with thermal-fluid systems. As a result, students may view thermodynamics as abstract, mathematically demanding, and disconnected from their disciplinary interests. Common learning challenges include defining appropriate systems and control volumes, interpreting property tables, connecting equations to physical processes, and choosing appropriate assumptions during problem solving.

Prior work in engineering education indicates that active learning, timely feedback, and structured opportunities for self-regulated learning can improve student engagement and performance in demanding STEM courses [1]-[6]. However, access to individualized support is often limited by staffing, office-hour availability, and the scale of multi-section service courses. Generative AI systems create new possibilities for just-in-time instructional support, but their use in engineering education must be intentionally designed, ethically implemented, and empirically evaluated.

Prior studies on educational chatbots and conversational agents suggest that these tools may support active learning, student engagement, and on-demand feedback when they are intentionally integrated into the learning environment rather than used as generic answer generators [2], [3], [7].

This study investigates a GPT-powered AI assistant as a supplemental instructional resource for Engineering Thermodynamics. The assistant is intended to support self-regulated learning [1], [10] and cognitive apprenticeship [11] by prompting students to articulate assumptions, select governing principles, explain intermediate steps, and reflect on whether their answers are physically reasonable. The assistant is not intended to replace human instruction, teaching assistants, or peer collaboration. Instead, it is designed to extend access to structured guidance outside the moments when the instructor or teaching assistants are immediately available.

This project is supported by a $10,000 Promoting Innovative Education Research (PIER) grant at the author's institution and is part of a broader research agenda on responsible AI integration in engineering education. The WIP contribution of this paper is the detailed design and study framework for deploying an AI assistant in a thermodynamics course, including safeguards for privacy, academic integrity, and interpretation of results.

The instructional challenge is not only conceptual but also logistical. Students often need feedback at the moment they are attempting to interpret a problem, choose assumptions, consult property tables, or decide which law applies. Traditional support mechanisms, including office hours, recitation sessions, and teaching assistants, are essential but limited by schedules, staffing, and student willingness to seek help. Generative AI tools create new possibilities for responsive instructional support, but they also introduce risks related to accuracy, overreliance, academic integrity, privacy, and uneven student use.

This paper presents a Work-in-Progress design framework for an AI-based thermodynamics teaching assistant. The initial project plan included deployment during the current semester. However, the tutor was not deployed because development was not complete and IRB approval for student participation and data collection was still pending. This delay revealed a central lesson of this work: in a required engineering course, an AI tutor must be treated as an instructional intervention, not merely as a software tool. It requires careful alignment with learning objectives, validation against course content, transparency about appropriate use, and safeguards that encourage learning rather than shortcut completion.

The purpose of this WIP paper is therefore to report the development status, design logic, IRB-dependent implementation barriers, and revised study plan. The contribution is not a completed efficacy study. Rather, it is a practical and theoretically grounded account of how an AI tutor for thermodynamics can be designed, evaluated, and prepared for future classroom deployment.

The future classroom study will be guided by the following research questions:

1. *To what extent does a GPT-based AI assistant improve students' conceptual understanding of thermodynamic principles compared to traditional instruction?*
2. *How does access to an AI assistant influence students' confidence and engagement, particularly among multidisciplinary students?*
3. *How do students interact with the assistant during problem-solving, and what types of support do they seek?*

The proposed AI assistant is grounded in self-regulated learning and cognitive apprenticeship. From a self-regulated learning perspective, the tutor is intended to help students monitor their understanding, identify gaps, request clarification, and practice problem-solving strategies. From a cognitive-apprenticeship perspective, the tutor is designed to make expert reasoning more visible by modeling how to define a system, select governing principles, articulate assumptions, and check whether an answer is physically reasonable.

The project also draws on active-learning and engagement frameworks that emphasize the value of feedback, explanation, and constructive student activity. Rather than positioning the AI assistant as a replacement for instruction, the design positions it as a scaffold for independent problem-solving. The assistant is intended to support students as they reason through problems, not to provide unexamined final answers.

The revised design no longer frames the assistant as adapting to fixed student learning styles. Instead, it emphasizes multiple representations and instructional supports. This revision also responds to longstanding discussions in engineering education about variation in how students engage with explanations, examples, and representations, while avoiding the stronger claim that students should be categorized into fixed learning-style groups [12]. For example, a student may receive a conceptual explanation, a step-by-step derivation, a table-interpretation guide, or a system diagram prompt depending on the nature of the question and the learning objective. This framing avoids assuming stable learner categories while still recognizing that students benefit from varied modes of explanation.

AI Assistant Design and Development

The AI assistant is being developed as a course-specific support tool for Engineering Thermodynamics. The intended use cases include clarifying conceptual questions, guiding students through problem setup, helping students interpret thermodynamic property data, prompting students to check units and assumptions, and suggesting next steps when students are stuck. The assistant is not intended to complete graded work for students or replace instructor feedback.

The prototype architecture includes a student-facing interface, a prompt-template layer, a course-content knowledge base, a LangChain-based orchestration layer, and integration with a large-language-model API. The student interface is planned to be implemented through a lightweight web application. The course-content knowledge base will include instructor-created materials, selected worked examples, course policies, and common misconception prompts. The selection and organization of course content will be guided by backward-design principles, beginning with course learning objectives and then aligning prompts, examples, and assessment plans with the desired evidence of student understanding [14]. The prompt templates will constrain the assistant to emphasize explanation, questioning, and reasoning steps rather than direct answer generation. This design approach is informed by human–AI interaction guidelines that emphasize transparency, controllability, and appropriate reliance, as well as prior work on designing complementary roles for teachers and AI systems in classroom learning environments [8], [9].

Figure 1 outlines the technical architecture of the assistant. The system allows dynamic prompt construction based on user input, instructional goals, and context-aware feedback.

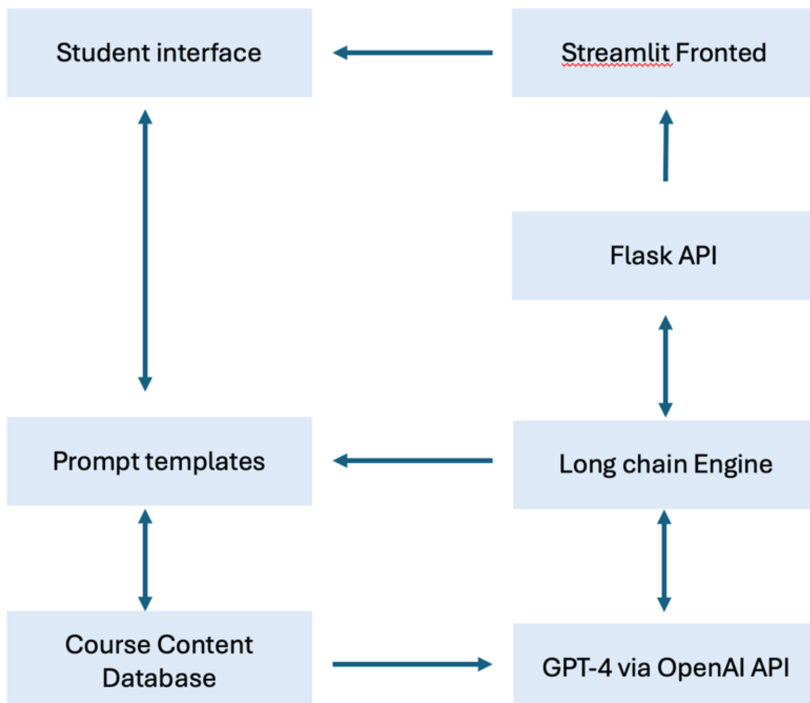

Figure 1. AI Assistant Architecture

Figure 2 provides a conceptual illustration of the same layered design, emphasizing how students interact with a constrained instructional interface rather than directly with an unconstrained language model.

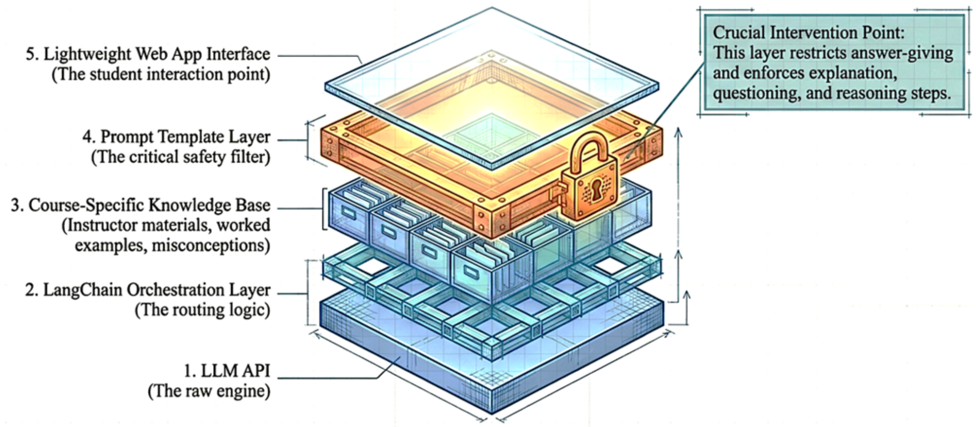

Figure 2. Conceptual layered architecture illustrating how the student-facing interface is separated from the raw large-language-model engine.

Three categories of prompt behavior are planned. Conceptual prompts will ask students to explain principles in their own words and connect equations to physical meaning. Problem-solving prompts will guide students through system definition, known and unknown variables, governing equations, assumptions, and solution checks. Metacognitive prompts will ask students to reflect on what they tried, why a selected method is appropriate, and whether the result is reasonable.

The development process to date has shown that course-specific validation is a substantial task. Thermodynamics questions often depend on diagrams, property tables, state assumptions, and problem context. A generic chatbot may provide plausible but incomplete or inappropriate guidance. For this reason, the assistant must be tested against representative course problems before classroom use. The validation process will include instructor review of responses, checks against known solutions, and refinement of prompt templates to reduce direct answer-giving and hallucinated assumptions.

Implementation delay and lessons learned

The original project plan was to deploy the AI tutor during the current semester and compare student outcomes across two sections. This deployment did not occur because the tutor was not sufficiently developed and validated for responsible use with students and because IRB approval for student testing and data collection was still pending. No students used the tutor as part of a research study, and no student learning data, usage data, or perception data were collected. While this prevented outcome analysis, it produced several useful implementation lessons for future AI-in-education work.

First, technical development required more time than anticipated. Building a course-specific AI tutor involves more than connecting a chatbot interface to a large-language-model API. It requires prompt design, response constraints, course-content organization, logging decisions, instructor-facing review mechanisms, and a usable interface that students can access without unnecessary friction.

Second, content validity is critical in thermodynamics. Even when a model produces coherent explanations, it may make hidden assumptions, skip important state definitions, or select an inappropriate property relation. For a required engineering course, these risks must be addressed before deployment through structured testing and instructor review.

Third, responsible use and academic integrity must be designed into the tool from the beginning. The tutor should promote reasoning, not automate homework completion. The revised design therefore includes guardrails that ask students to identify the problem context, state assumptions, and explain their reasoning before receiving detailed guidance. For graded assignments, the assistant can be limited to conceptual hints, process guidance, and feedback on student-generated steps rather than complete solutions.

Fourth, ethical and privacy considerations must be resolved before any student interaction data are collected. At the time of this WIP paper, IRB approval is still pending; therefore, student testing has not begun. The future implementation will proceed only after approval is obtained and will use informed consent, de-identified records, secure storage, and clear communication about what data are collected and how they will be used. Student participation in research data collection will be voluntary and separated from course grading to the extent possible under institutional review requirements.

These lessons reframed the project from a simple classroom intervention to a responsible-design and approval-dependent deployment problem. The delayed implementation is therefore not treated as a failed result, but as evidence that careful preparation is necessary before introducing generative AI into a high-stakes learning environment.

Methodology

After IRB approval and tool validation are complete, the planned future study will use a mixed-methods design. Two sections of Engineering Thermodynamics will be compared, with one section using the AI assistant in supervised or clearly bounded settings and the other section receiving standard instruction. If the study is approved and implemented, both sections will receive equivalent lectures, assignments, assessments, and access to regular instructional support. The same instructor will teach both sections when possible; if multiple teaching assistants are involved, TA responsibilities, recitation materials, and grading rubrics will be standardized to reduce instructor and TA effects.

Equivalence between sections will be examined in the future study at baseline using available non-sensitive academic and demographic indicators permitted under institutional review protocols. These may include major, academic level, prior completion of prerequisite coursework, baseline concept-inventory score, and self-reported confidence with thermodynamics-related problem solving. These baseline measures will be used descriptively and, where appropriate, as covariates or matching variables in the analysis. Table 1 summarizes the planned comparison between the control and AI-assisted sections.

Table 1. Group A vs. Group B information

Component	Comparison section	AI-assistant section
Lecture format	Standard course lectures	Same standard course lectures

Component	Comparison section	AI-assistant section
Recitation/support	Standard recitation and TA support	Standard support plus bounded AI-assistant access
Course assessments	Same quizzes, exams, and assignments	Same quizzes, exams, and assignments
Baseline measures	Concept inventory, confidence survey, background indicators	Concept inventory, confidence survey, background indicators
Usage data	Not applicable	De-identified AI usage logs, if consented
Ethics/consent	IRB approval pending; consent procedures planned	IRB approval pending; consent procedures planned

The planned concept inventory will focus on central thermodynamics topics emphasized in the course, including system boundaries, the first law, property relationships, entropy, and the second law. If an existing validated thermodynamics concept inventory is available and aligned with the course outcomes, selected items will be adapted with permission or used as a model. If a locally developed instrument is required, items will be reviewed by thermodynamics instructors for content validity and piloted before being used for formal analysis. The paper will clearly distinguish between validated, adapted, and instructor-developed items in any future reporting. The revised study design maps each research question to specific data sources, as shown in Table 2.

Table 2. Research-question alignment with data sources and planned analyses.

Research question	Primary data sources	Planned analysis
RQ1: Conceptual understanding	Pre/post concept inventory, common quiz and exam scores	Paired t-tests for within-group gains; independent-sample t-tests or ANCOVA-style exploratory comparisons when appropriate
RQ2: Confidence and engagement	Surveys, homework engagement, attendance or participation indicators	Descriptive statistics, group comparisons, and thematic analysis of open-ended responses
RQ3: Interaction patterns	AI usage logs, query types, duration/frequency, TA observation notes	Usage categorization, correlation analyses, and qualitative coding

Table 3. Planned future study design overview after IRB approval.

Week(s)	Activity	AI use mode	Data collected
1-2	Baseline orientation and concept check	No AI use	Concept inventory pre-test; AI readiness survey; consent documentation

Week(s)	Activity	AI use mode	Data collected
3-4	Introduction to appropriate AI use	Demonstration only	Student questions; instructor and TA notes
5-12	AI-enabled recitations and guided practice	Supervised use in treatment section	Planned usage logs, reflections, TA notes, and common-misconception tracking, pending IRB approval.
13	Post-intervention concept check	No AI during assessment	Concept inventory post-test; engagement survey
14-15	Interviews or focus groups	Not applicable	Semi-structured transcripts and open-ended feedback

Data collection and analysis
After approval is obtained, quantitative data will include pre- and post-concept-inventory scores, quiz and exam performance, homework or optional-practice engagement, and AI usage metrics for consenting students in the treatment section. Usage metrics may include number of sessions, frequency of use, approximate duration, broad query category, and whether students request conceptual clarification, procedural guidance, or answer-checking support. No personally identifying interaction content will be reported.

After approval is obtained, qualitative data will include open-ended survey responses, optional focus group comments, instructor reflections, and TA observation notes. These data will be used to understand how students perceive the tool, what barriers they encounter, and whether they view the assistant as supporting learning or encouraging overreliance. Qualitative responses will be coded thematically by at least two reviewers when feasible, with disagreements resolved through discussion.

Because classroom deployment has not occurred and IRB approval is still pending, no student outcome data, usage data, or student-perception data are presented in this WIP paper. Future analysis will compare within-group pre/post changes and between-group differences while accounting for baseline differences where possible. The analysis will also examine whether patterns of AI use are associated with learning gains, confidence, or engagement. These analyses will be interpreted cautiously because students self-select how frequently they use optional support tools.

Academic integrity and responsible-use safeguards
A central design concern is ensuring that the AI assistant supports learning without becoming a shortcut for graded work. The planned assistant will include several safeguards. First, the interface and course policy will communicate that the tool is for learning support, not answer outsourcing. Second, prompt templates will prioritize Socratic questioning, partial hints, and requests for student reasoning before providing detailed guidance. Third, for homework-like

questions, the assistant will be instructed to avoid giving final numerical answers unless the student has already shown a complete solution attempt or the interaction occurs in a supervised practice setting.

Additional safeguards may include instructor review of anonymized query patterns, explicit examples of acceptable and unacceptable AI use, and alignment with the course academic-integrity policy. The goal is to help students develop judgment about when and how AI can support engineering problem solving, while maintaining responsibility for their own work. This approach is consistent with broader calls for AI in education to be implemented as a guided, transparent, and pedagogically purposeful support system rather than as an unregulated replacement for student reasoning or instructor judgment [13].

Anticipated outcomes and contribution

The future study is expected to clarify whether a course-specific AI assistant can improve student access to timely feedback and support self-regulated learning in thermodynamics. However, this WIP paper does not claim that the assistant improves performance, engagement, or confidence because those outcomes have not yet been measured. Instead, the current contribution is a design and implementation framework for responsible deployment.

This contribution is relevant to engineering educators because many AI-in-education projects face a similar tension: generative AI tools are easy to access but difficult to integrate responsibly into required technical courses. The development experience reported here suggests that educators should allocate time for course-specific validation, academic-integrity planning, student-data protections, and instructor oversight before deployment. These elements are not secondary administrative concerns; they are central to the instructional design of AI-supported learning environments.

The revised implementation plan will inform the next phase of the project. Once the tutor has been validated and IRB approval has been obtained, the tool may be deployed in a future course offering and student learning and engagement data can be collected and reported in a subsequent study.

The primary limitation of this WIP paper is that the AI assistant was not deployed to students during the originally planned semester because development was still underway and IRB approval was pending. As a result, the paper does not include student learning outcomes, usage data, or student perceptions from an implemented intervention. The paper is therefore best understood as a design and implementation-planning study rather than an effectiveness study.

A second limitation is that the planned future comparison between course sections may be affected by non-random section enrollment, differences in student schedules, and variation in prior preparation. The revised study will address these issues through baseline measures and careful reporting, but causal claims will remain limited unless random assignment or stronger quasi-experimental controls become feasible.

Conclusion

This Work-in-Progress paper presents the design status, IRB-pending deployment, and revised implementation framework for an AI-based teaching assistant in Engineering Thermodynamics. The inability to deploy the tutor during the originally planned semester shifted the project from an outcome-focused pilot to a responsible-design and planning study. This shift is important because generative AI tools require careful validation before being introduced into a required engineering course.

The revised framework emphasizes course-specific content alignment, multiple representations of thermodynamic reasoning, bounded student use, academic-integrity safeguards, privacy protections, and a mixed-methods evaluation plan. Future work will complete validation of the assistant, obtain IRB approval, deploy the tool in a controlled course setting, and examine its effects on conceptual understanding, confidence, engagement, and help-seeking behavior. By documenting both the promise and the implementation constraints of AI-assisted learning, this paper contributes practical guidance for educators seeking to integrate generative AI into engineering education responsibly.

References:

[1] A. C. Graesser and D. S. McNamara, "Self-regulated learning in learning environments with pedagogical agents that interact in natural language," *Educational Psychologist,* 2010.

[2] R. Winkler and D. Soergel, "Conversational Agents in Education: Supporting Active Learning in STEM," *Computers & Education,* 2016

[3] D. E. Gonda and A. H. Kerly, "Application of Chatbot for Computer Science Learning: A Review of Benefits and Challenges," *Journal of Education and e-Learning Research,* 2021.

[4] K. A. Case and G. Light, "Emerging methods for teaching thermodynamics: Engaging students in abstract concepts," *Journal of Engineering Education*, vol. 100, no. 4, pp. 723–732, 2011.

[5] M. Prince and R. M. Felder, "Inductive teaching and learning methods: Definitions, comparisons, and research bases," *Journal of Engineering Education*, vol. 95, no. 2, pp. 123–138, 2006.

[6] M. T. H. Chi, M. Roy, and R. G. Wylie, "The ICAP framework: Linking cognitive engagement to active learning outcomes," *Educational Psychologist*, vol. 49, no. 4, pp. 219–243, 2014.

[7] R. Winkler and M. Söllner, "Unleashing the potential of chatbots in education: A state-of-the-art analysis," *Computers & Education*, vol. 120, pp. 1–14, 2018.

[8] S. Amershi *et al*., "Guidelines for human–AI interaction," in *Proc. CHI Conference on Human Factors in Computing Systems*, Glasgow, UK, 2019.

[9] K. Holstein, G. McLaren, and V. Aleven, "Designing for complementarity: Teacher and AI roles in orchestration of classroom learning," in *Proc. International Conference on Artificial Intelligence in Education (AIED)*, 2020.

[10] B. J. Zimmerman, "Becoming a self-regulated learner: An overview," *Theory Into Practice*, vol. 41, no. 2, pp. 64–70, 2002.

[11] A. Collins, J. S. Brown, and S. E. Newman, "Cognitive apprenticeship: Teaching the craft of reading, writing, and mathematics," in *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*, Hillsdale, NJ, USA: Lawrence Erlbaum Associates, 1989, pp. 453–494.

[12] R. M. Felder and L. K. Silverman, “Learning and teaching styles in engineering education,” *Engineering Education*, vol. 78, no. 7, pp. 674–681, 1988.
[13] R. Luckin, W. Holmes, M. Griffiths, and L. B. Forcier, *Intelligence Unleashed: An Argument for AI in Education*. London, UK: Pearson Education, 2016.
[14] G. Wiggins and J. McTighe, *Understanding by Design*, 2nd ed. Alexandria, VA, USA: ASCD, 2005.