Faculty Research

Enhancing the Academic Advising Process with Large Language Models

This white paper provides a summary of the work from:

Kuzlu M, Jovanovic VM, Smith K, Popescu O, El-Shahat A, Cotton R, Seegers KM, Henderson WA. A Large Language Model-Based Academic Advising Assistant for Engineering Technology Students. ASEE Annual Conference & Exposition, 2026.

Research Author Bios:

Murat Kuzlu, Associate Professor; Vukica M. Jovanovic, Professor and Department Chair; Katherine Smith, Assistant Professor; Otilia Popescu, Professor; Adel El-Shahat, Assistant Professor; Ryan Cotton, Undergraduate Student; Kayla Marie Seegers, Undergraduate Student; and William Austin Henderson, Undergraduate Student - Old Dominion University

White Paper Author:

Murat Kuzlu, Ph.D., is an Associate Professor in the Department of Engineering Technology at Old Dominion University (ODU).

Overview

Academic advising plays a crucial role, particularly for undergraduate students, in making continuous and controlled progress toward graduation. However, in recent years, academic advising has become more complicated, as course pre- or corequisites and availability, department and program requirements, and individual circumstances must all be considered at the same time. Large language models (LLMs) have great potential to make this process more effective and interactive by allowing students to describe their specific needs and providing a multi-semester plan with explanations rather than only lists of remaining courses and other requirements. This white paper summarizes an LLM-based academic advising assistant system developed for Engineering Technology students at ODU.

Purpose of the Research

The main objective of this study is to develop an LLM-based academic advising system, utilizing a student's academic performance, preferences, and constraints, as well as university and program requirements, to recommend a personalized multi-semester course plan. The study also aims to answer the following research question:

- How can an LLM-based academic advising system contribute to improving the current academic advising process without replacing human academic advisors?

Method

Figure 1 shows how the proposed academic advising assistant works in three steps, i.e., steps 1 and 2: obtaining advisor and student input, and step 3: combining the obtained information and providing it to the selected LLM. In

this process, first, advisors enter information about degree programs, courses, and other requirements. Second, students select their program, identify courses they have completed, and specify preferences and constraints, such as how many courses they want to take each semester or when they plan to graduate. Finally, the backend system uses this information to generate a personalized course plan for upcoming semesters and explain why the model recommends this plan.

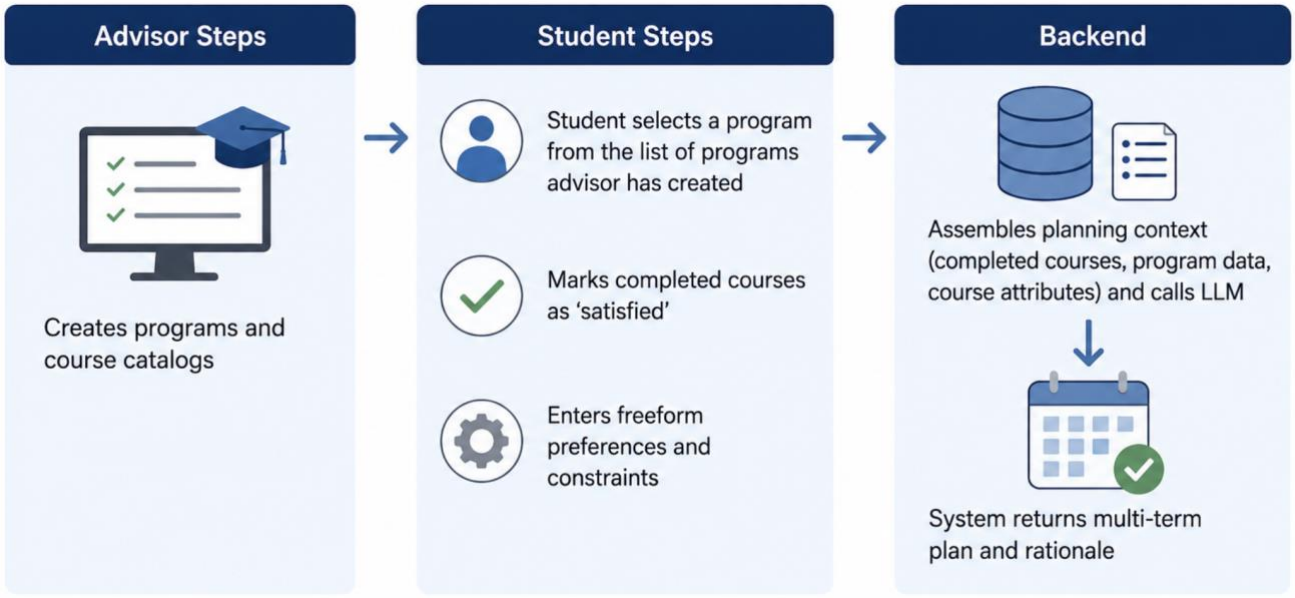


Figure 1. The proposed LLM-based academic advising workflow.

Figure 2 provides a more detailed view of the technical architecture behind the developed academic advising system. In the system, the web-based user interface presents the front end, providing an interactive dashboard to collect relevant information from students and advisors. The database (PostgreSQL) stores that information and other academic requirements, and an AI model (LLM) utilizes all information to generate a multi-semester plan. Students select their academic program, catalog year, and completed courses, choose a starting semester, and enter scheduling preferences along with constraints via the dashboard. Students can also limit the number of courses or credits in a semester or ask to avoid particular course combinations based on their preferences. The academic advising system then combines this information with academic data stored in the database, including programs, program concentrations, course credits, prerequisites, and availability.

Figures 3 and 4 show examples of the student and advisor dashboards, respectively. The team built the dashboard using Streamlit, a Python-based framework for interactive applications. When a student requests an academic plan, the application gathers the relevant course information and their preferences and sends them to the selected LLM, i.e., OpenAI's GPT-4o-mini. The model organizes courses and generates a multi-semester plan based on course pre- or corequisites and availability, considers student preferences and constraints, balances the workload, and explains the proposed multi-semester plan. Students can view the generated academic plan in the system and save it for later review.

The team also plans to add a set of automated checks for mandatory university requirements, such as prerequisites and other degree rules. In the current prototype, many of these requirements are provided to the LLM through instructions.

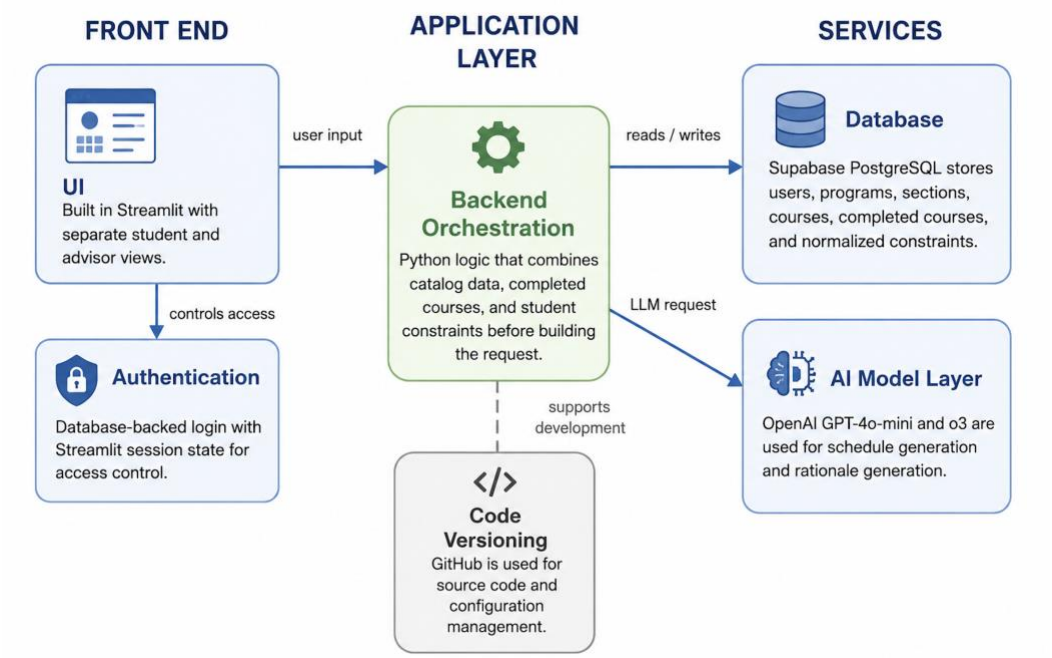
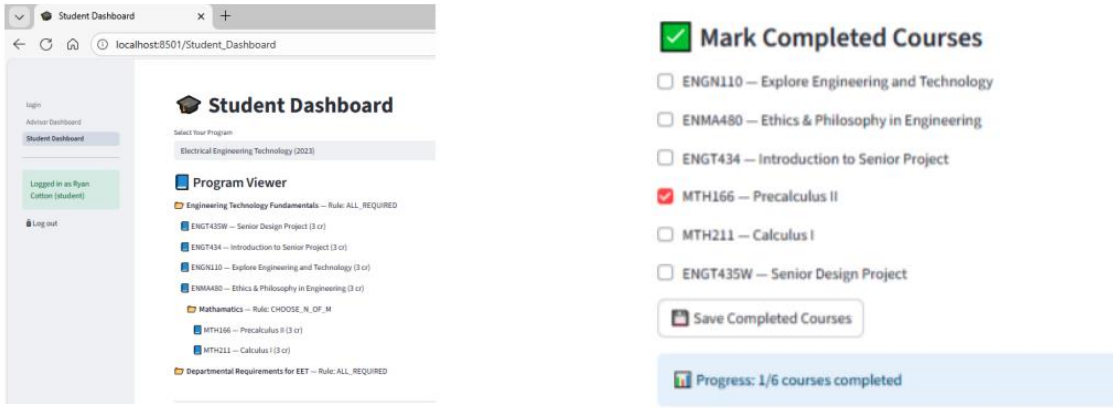


Figure 2. The technical architecture of the academic advising assistant and its implementation stack.



(a) (b)

Figure 3. View of the student dashboard: (a) Program Viewer and (b) Completed Courses.

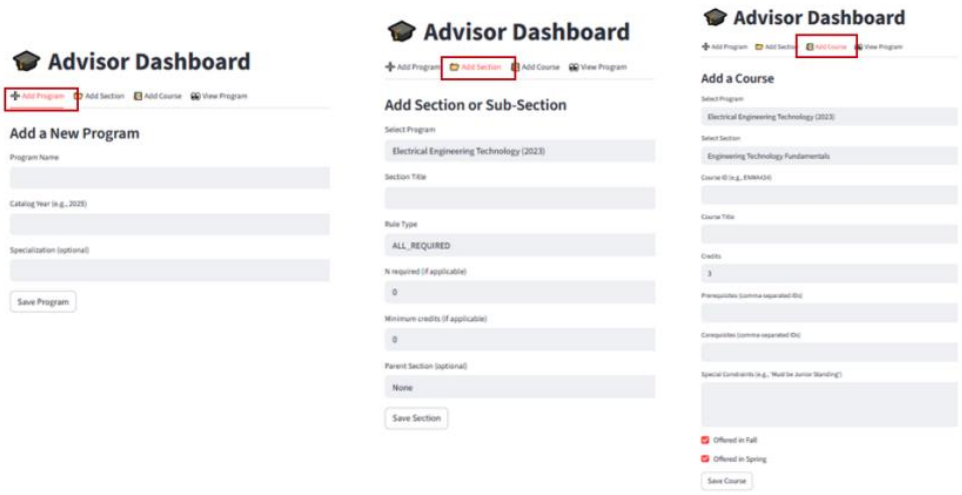


Figure 4. View of the advisor dashboard.

Findings and Discussion

The developed prototype demonstrates how an LLM-based advising experience can move beyond the current static advising systems. In this system, a student can easily review the program requirements, identify courses (completed and to be completed), and then the AI Planner provides the multi-semester plan to students based on this input and explains the reasoning behind the recommendation. The combination of planning with explanation is the key feature of the developed advising system for students, i.e., helping students understand not only what remains in a degree, but also why courses should be taken in a particular sequence. Examples of the AI Planner and the AI-generated multi-semester plan are given in Figure 5.

Another key feature of the academic advising system is personalization. The system is designed to generate an academic plan for each student based on their own academic history and constraints rather than generating one standard schedule for everyone. For example, a student may want to limit the number of courses per semester due to other commitments or avoid taking two demanding courses in the same semester. The AI Planner incorporates these preferences when generating the multi-semester schedule for a student. This offers a unique opportunity for students to explore alternatives before meeting with an advisor. The study also highlights an important lesson learned for universities considering generative AI for advising, i.e., LLMs should not be the sole source of academic rules. In addition, LLMs can generate incorrect information or make unsupported assumptions when they are used without sufficient structure. This is a serious trustworthiness concern for the academic advising process because a mistaken prerequisite or course offering could cause a delay in graduation. This also makes verified data and deterministic validation essential. Therefore, the current system should be considered a prototype and demonstration rather than a self-sufficient advising service.

For faculty and advisors, LLMs are probably best as an extra layer of support for advising systems, not a replacement. An LLM-based academic advising assistant system could answer routine questions related to program and course requirements and scheduling, help students visualize possible multi-semester plans, and explain them with course dependencies on demand. In the system, human advisors remain central to the academic advising process and are responsible for final decisions and exceptions. In this way, LLM-based solutions may help make the advising process more accessible while preserving the judgment and personal support provided by faculty and professional advisors.

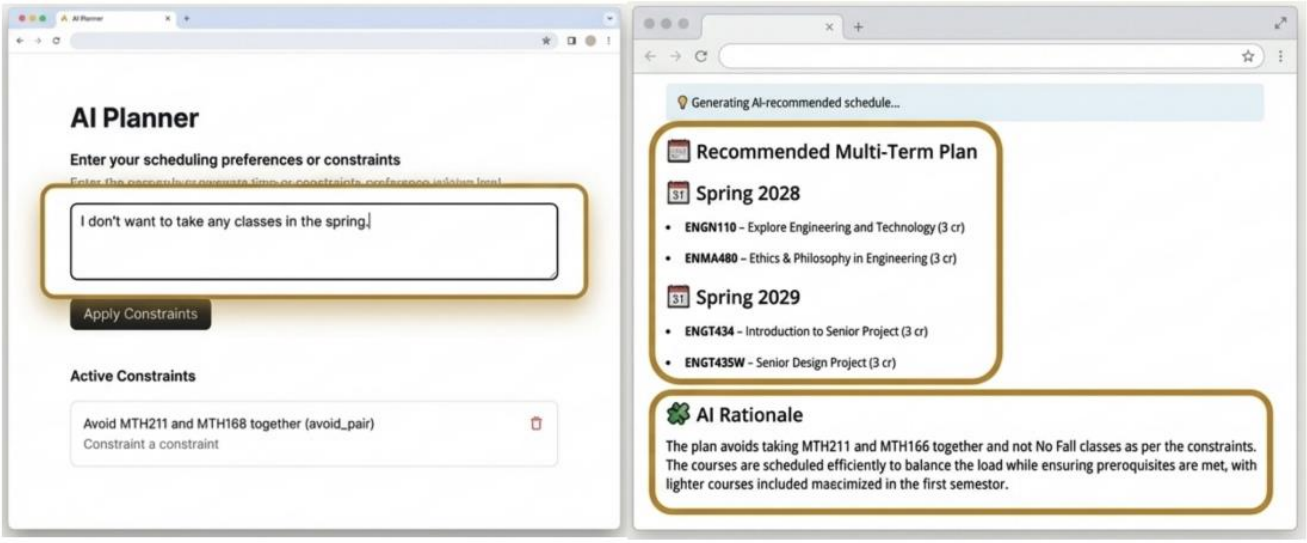


Figure 5. (a) AI Planner and (b) AI-generated multi-semester plan.

Conclusions

Many higher education institutions have observed an increasing demand for academic advising services in degree planning. However, many current advising systems struggle to support complex and multi-step reasoning. Large

language models (LLMs) offer an opportunity to make academic advising more interactive and easier for both students and advisors to navigate. The main objective of the proposed academic advising system is to address this issue and utilize LLMs to provide an additional on-demand planning resource to students while providing a tool that can help with repetitive planning tasks for advisors. The developed prototype shows how natural-language interaction can be combined with students' inputs and constraints, course information, program requirements, and scheduling preferences to create a personalized multi-semester plan.

This white paper also emphasizes the importance of an LLM-based advising assistant by demonstrating knowledge of contextual understanding, explainability, and multi-step reasoning, presenting an opportunity to improve academic advising systems, particularly for higher education institutions. Findings indicate that verified institutional data, application logic, deterministic controls, secure system design, and advisor oversight are all important parts of the overall solution.

References

Kuzlu M, Jovanovic VM, Smith K, Popescu O, El-Shahat A, Cotton R, Seegers KM, Henderson WA. A Large Language Model-Based Academic Advising Assistant for Engineering Technology Students. ASEE Annual Conference & Exposition, 2026.