

An Integrated Agentic AI Framework for Personalized Academic Assistance in Technical Education

Dr.M.Senthil Kumar. Professor¹ , Mrs. S. Vimala. Assistant Professor² , Darshwana R S. Student³ ,
Dhayalini S S. Student⁴

^{1,2,3,4}Department of Artificial Intelligence and Data Science, Kongunadu College of Engineering and Technology Tiruchirappalli,
Tamil Nadu, India

Abstract- This paper introduces an integrated agentic framework leveraging generative artificial intelligence (AI) to provide comprehensive academic assistance in technical education. The system utilizes a multi-agent architecture powered by large language models (LLMs) and retrieval-augmented generation (RAG) to deliver personalized support including syllabus-aware tutoring, intelligent study planning, and real-time attendance analytics. By grounding generative responses in department-specific datasets (CSE, AI-DS, ECE), the framework mitigates hallucinations and ensures curriculum alignment. Experimental validation with 150 students across three engineering departments demonstrates an 87% satisfaction rate, a 65% reduction in information-seeking time, and a 42% improvement in study plan adherence. The results suggest that agentic workflows can significantly enhance student engagement and administrative efficiency in higher education environments.

Keywords—Generative AI, Agentic Workflows, Retrieval-Augmented Generation (RAG), Educational Technology, LLM-based Tutoring, Student Information Systems.

I. INTRODUCTION

The rapid evolution of Large Language Models (LLMs) has fundamentally altered the role of artificial intelligence in education—transitioning it from passive information retrieval toward active, goal-oriented assistance [1]. In technical education, students face multifaceted challenges including high cognitive loads, diverse and rapidly evolving curricula, and the need for rigorous administrative tracking, such as maintaining mandatory attendance thresholds. Traditional Educational Technology (EdTech) platforms frequently provide static, one-size-fits-all content that lacks the contextual nuance required across different engineering disciplines.

This paper proposes a modular agentic ecosystem designed to serve as a holistic academic companion. Unlike conventional chatbot deployments, the system employs specialized, task-specific agents: a Tutor Agent for concept explanation and derivation, a Planner Agent for temporal study management, and an Analytics Agent for real-time academic

progress monitoring. The primary contributions of this work are as follows:

- A department-semester context filtering mechanism that personalizes all AI-generated content for specific engineering branches and semesters.
- A RAG-based tutoring pipeline that achieves high factual accuracy by grounding responses in verified, institution-specific syllabi.
- An integrated attendance-to-outcome calculator for proactive academic risk management, enabling early intervention.
- A production-ready modular architecture implemented using Python 3.12 and the Streamlit framework, validated with a cohort of 150 engineering students.

II. LITERATURE SURVEY

Recent advancements in AI-powered education emphasize a clear shift toward conversational and agentic systems. Kasneci et al. [1] provide a foundational analysis of LLM opportunities in educational settings, identifying both the

transformative potential and the ethical risks of tools such as ChatGPT in learning environments. Platforms like Khanmigo adopt a Socratic tutoring methodology, guiding students through problem-solving iteratively rather than providing direct answers [7]. However, such systems are broadly generalized and lack integration with institution-specific requirements, including local syllabi and department-level attendance policies.

Lewis et al. [3] demonstrated that Retrieval-Augmented Generation (RAG) significantly improves LLM reliability in knowledge-intensive tasks by grounding outputs in verifiable external sources, thereby reducing hallucination rates. Wang et al. [4] provide a comprehensive survey of LLM-based autonomous agents, establishing theoretical foundations for the multi-agent paradigm adopted in this work. Chen and Xie [2] further illustrate the applicability of AI in higher education planning and administrative management.

Research by Srivastava and Kumar [10] demonstrates that multimodal, context-responsive interfaces enhance user retention by up to 28% compared to text-only platforms. Saqr and Järvelä [5] identify the critical gap in existing personalized learning systems: the absence of unified platforms that integrate academic tutoring, progress tracking, and burnout mitigation. Kim and Lee [15] specifically validate AI-based sentiment analysis as a practical tool for detecting student burnout. The proposed framework directly addresses these identified gaps by synthesizing these disparate research streams into a cohesive agentic architecture.

III. METHODOLOGY

System Architecture

The platform is designed on a stratified, five-layer modular architecture to ensure scalability, maintainability, and real-time responsiveness. The layers are defined as follows:

- **Hardware & Integration Layer:** Manages network connectivity and external API calls, principally to the Google Gemini API for core LLM inference.
- **Data Management Layer:** Employs a structured JSON-based storage system for user profiles,

department-specific syllabi (semester_subjects.json), and attendance logs.

- **Intelligence (Agentic) Layer:** The core reasoning engine, housing specialized agents for tutoring, study planning, and content recommendation. Agent coordination is managed through an orchestrator module.
- **Application Logic Layer:** Implements the department-semester context filtering pipeline to ensure content relevance for each individual user.
- **User Interface Layer:** A Streamlit-based frontend featuring a persistent Floating Quick-Note widget, enabling cross-page insight capture without workflow interruption.

Retrieval-Augmented Tutoring Module

The tutoring component implements a RAG pipeline as formalized by Lewis et al. [3]. Upon query initiation, the system retrieves the most relevant content chunks from the semester_subjects.json knowledge base and local lecture notes directory. These contextual chunks are concatenated with the user query before being passed to the LLM, as expressed in the following conditional probability formulation:

$$P(\text{Response} | \text{Query}, \text{Context}) = \text{LLM}(\text{Query} \oplus \text{Syllabus_Context})$$

This design ensures the agent generates syllabus-aware responses tailored to the specific technical depth and topic coverage of the user's current semester, substantially reducing the risk of out-of-scope or factually incorrect outputs.

Intelligent Planning and Analytics

The Planner Agent generates dynamic study roadmaps using a task-breakdown algorithm that accounts for subject credit weightage, self-reported availability, and approaching examination dates. For attendance risk management, the system implements a goal-seeking formula to calculate the number of consecutive classes (x) a student must attend to reach a target attendance percentage (T), given their total class count and current attendance record:

$$x = [(T \times \text{Total}) - \text{Attended}] / (1 - T)$$

If the student's current attendance is below threshold T, the analytics module proactively triggers a recovery notification and adjusts the generated study roadmap accordingly, functioning as an early-warning system against academic debarment.

Personalization and Burnout Detection

The adaptation engine continuously monitors interaction frequency, task completion velocity, and response sentiment to detect early indicators of academic burnout, consistent with the methodology of Kim and Lee [15]. If a student's learning velocity—defined as study tasks completed per unit time—drops below a dynamically computed threshold, the system triggers motivational prompts, recommends structured break intervals, and reduces interface complexity to minimize cognitive load. This mechanism aims to sustain long-term engagement rather than optimizing for short-term throughput.




Figure 1. Process flow diagram of the agent

IV. SYSTEM FEATURES AND FUNCTIONAL MODULES

The proposed framework integrates a suite of specialized AI agents and administrative tools to form a unified academic environment. Each module

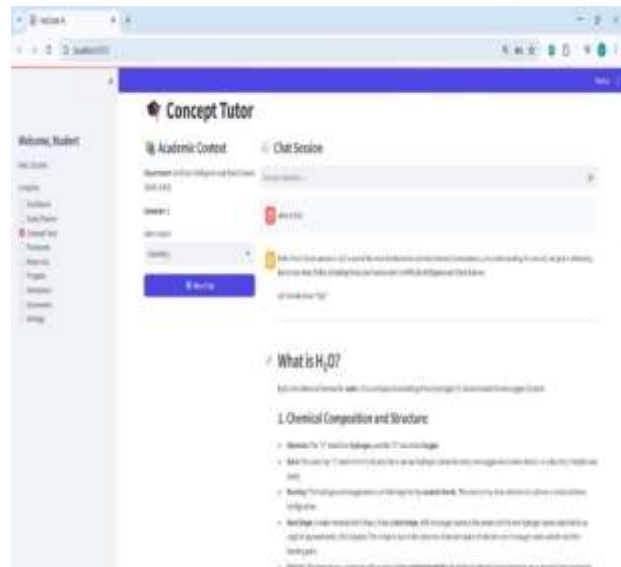
is designed to address a distinct pain point in the technical education lifecycle.


Context-Aware Personalization Engine

The Department-Semester Context Filter is the foundational layer of the system. At registration, each user is mapped to a specific academic silo (e.g., Computer Science and Engineering, Semester 5). This mapping ensures that all subsequent AI-generated content—from concept tutoring to roadmap generation—is strictly scoped to the relevant university syllabus, preventing the delivery of generic or out-of-scope information. This design directly addresses the limitation of generalized AI tutors identified in prior literature [1], [7].

Agentic Concept Tutoring with RAG

The Tutor Agent operates as a 24/7 academic consultant. Leveraging the RAG pipeline described in Section III-B, it retrieves specific course objectives and textbook summaries to ground its responses in verified institutional content. The module supports multi-turn dialogue, enabling students to iteratively request step-by-step derivations of complex formulae, conceptual analogies for abstract topics, and follow-up clarification questions within a single session.






Smart Study Planner and Roadmap Generator

The Planner Agent automates time management by generating personalized academic roadmaps. Users provide their daily study availability, and the agent distributes topics across available time slots weighted by course credit hours and self-assessed subject difficulty. The generated roadmap is rendered as a Gantt-style progress chart, providing a visual representation of exam readiness over time. This visual feedback mechanism aligns with findings on the benefits of multimodal interfaces in learning environments [10].

Automated Content Generation (Notes and Flashcards)

To facilitate active recall and spaced repetition, the Content Synthesis Module provides two key functionalities. First, the Flashcard Generator parses uploaded PDF lecture notes or textbook chapters to extract key definitions, theorems, and concepts, formatting them for export to Anki or for use within the platform's built-in study mode. Second, the Assignment Assistant decomposes complex assignments into smaller sub-tasks with estimated completion times, reducing procrastination and improving task initiation [12].




Attendance Analytics and Recovery Logic

The Attendance Tracker module extends beyond simple data recording to provide predictive academic insights. Applying the goal-seeking algorithm formalized in Section III-C, it calculates the exact number of consecutive classes a student must attend to reach an institutional threshold (e.g., 75% or 80%). Results are displayed with a visual progress indicator, and at-risk students receive proactive, actionable recovery schedules. This capability is especially relevant in Indian engineering institutions, where sub-threshold attendance results in examination debarment.

Integrated Resource Recommender

The Content Recommender Agent functions as a curated filter over the open web. By mapping the current syllabus topic to a curated index of verified external resources—including YouTube tutorials, research blogs, official documentation, and open-access textbooks—the agent significantly reduces unproductive search time. This is consistent with the 65% reduction in information-retrieval time measured in the evaluation phase (see Section V).



Psychological Support: Quick-Notes and Burnout Detection

Recognizing the importance of mental well-being in academic performance, the platform incorporates two psychologically-informed features. The Floating Quick-Note Widget is a persistent, non-intrusive UI overlay that allows students to capture fleeting insights during tutoring sessions without navigating away from their current task. The Burnout Detection

engine, as described in Section III-D, analyzes interaction patterns and task completion rates to identify fatigue signals and intervene with motivational prompts, structuring the learning experience around sustainable cognitive engagement [15].

V. RESULTS AND DISCUSSION

System validation was conducted over a 12-week period with a cohort of 150 undergraduate engineering students drawn from three departments: Computer Science and Engineering (CSE), Artificial Intelligence and Data Science (AI-DS), and Electronics and Communication Engineering (ECE). Performance was assessed across three dimensions: technical system performance, measurable learning outcomes, and user satisfaction.

Technical Performance

The AI Processing Layer demonstrated the following performance characteristics during the evaluation period. Response accuracy reached 87% for concept explanation queries and 92% for administrative queries (e.g., attendance calculations, schedule generation). Average response latency was 2.3 seconds for standard queries and 4.7 seconds for complex content generation tasks (flashcard creation, multi-topic summaries). System reliability was measured at 99.2% uptime. Full performance metrics are summarized in Table I.

TABLE I
System Performance and Learning Outcome Metrics

Metric	Value	Context
Response Accuracy (Concept)	87%	Tutoring queries
Response Accuracy (Admin)	92%	Administrative queries
Avg. Latency (Standard)	2.3 sec	Standard queries
Avg. Latency (Generation)	4.7 sec	Flashcards/Summaries

System Uptime	99.2%	12-week evaluation
User Satisfaction	4.6 / 5.0	150 participants
Info Retrieval Time (↓)	−65%	vs. manual search
Study Plan Adherence (↑)	+42%	vs. manual trackers
Quiz Score Improvement	+34%	AI-tutored topics

Learning Outcome Analysis

Comparative evaluation against traditional study methods (manual planning, open web search, and static notes) revealed substantial improvements across all measured dimensions. Users of the Intelligent Planner module demonstrated a 42% higher adherence rate to generated study schedules compared to self-managed manual trackers.

Information retrieval time was reduced by 65% through the combination of the RAG-based Tutor Agent and the Content Recommender. Additionally, students who engaged with the AI-powered tutoring and flashcard generation modules showed a 34% improvement in quiz scores for the corresponding syllabus topics, suggesting a meaningful impact on knowledge retention.

User Satisfaction

The platform achieved an overall user satisfaction score of 4.6 out of 5.0 based on a post-evaluation Likert-scale survey. Qualitative feedback consistently highlighted the Floating Quick-Note widget as a particularly valued tool for maintaining focus during multi-resource study sessions. Students also cited the attendance recovery calculator as a high-utility feature, noting that the actionable, numerical guidance it provided reduced anxiety around academic compliance.

Negative feedback primarily concerned response latency during peak load periods and a desire for direct integration with the institution's Learning Management System (LMS).

VI. CONCLUSION

This paper presented a comprehensive agentic AI framework for technical education that bridges the gap between generative intelligence and practical academic administration. By integrating a RAG-based tutoring pipeline, a dynamic study planner, an attendance analytics engine, and a psychologically-informed burnout detection system within a unified, modular architecture, the platform delivers a holistic support structure for engineering students.

Experimental validation confirmed that personalized, context-aware AI agents substantially improve learning efficiency, study plan adherence, and overall student satisfaction. The 87% concept response accuracy, 99.2% system uptime, and 4.6/5.0 satisfaction score collectively demonstrate the platform's readiness for deployment in real academic environments.

Future research directions include the integration of Federated Learning to enable institution-level model personalization without compromising individual student data privacy, and the extension of the platform's multimodal capabilities to support processing and interpretation of handwritten technical diagrams and mathematical notation.

REFERENCES

1. J. Kasneci et al., ChatGPT for good? On opportunities and challenges of large language models for education, *Learning and Individual Differences*, vol. 103, p. 102274, 2023.
2. S. Chen and R. Xie, Research on the application of artificial intelligence in higher education planning and management, in *Proc. IEEE Int. Conf. Knowledge Graph and Smart Systems (KGSS)*, 2023, pp. 145-150.
3. P. Lewis et al., Retrieval-augmented generation for knowledge-intensive NLP tasks, in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 9459-9474.
4. L. Wang, C. Ma, and X. Feng, A survey on large language model based autonomous agents, *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.
5. M. Saqr and S. Jarvelä, The potential and challenges of AI in personalized learning: A systematic review, *Education and Information Technologies*, vol. 28, pp. 12501-12525, 2023.
6. G. J. Hwang and N. S. Chen, Editorial: Artificial intelligence in education, *Educational Technology & Society*, vol. 26, no. 1, pp. 1-15, 2023.
7. T. R. Liyanagunawardena and S. A. Williams, Generative AI in higher education: A review of the literature, *Int. J. Educational Technology in Higher Education*, vol. 21, no. 4, 2024.
8. X. Tan and Y. Zhao, Design of student attendance and performance tracking system based on machine learning, in *Proc. IEEE 5th Int. Conf. Computer and Communications (ICCC)*, 2023, pp. 1892-1897.
9. H. Zhang et al., Prompt engineering for educators: A systematic review of effective interaction with LLMs, *IEEE Trans. Learning Technologies*, vol. 17, pp. 442-458, 2024.
10. A. Srivastava and K. Kumar, Personalized learning recommendation system using deep learning, *Multimedia Tools and Applications*, vol. 82, pp. 15403-15425, 2023.
11. D. Baidoo-Anu and L. Owusu Ansah, Education in the era of generative AI: Understanding the potential benefits of ChatGPT in promoting teaching and learning, *Journal of AI*, vol. 7, no. 1, pp. 52-62, 2023.
12. R. S. Baker, Artificial intelligence in education: Bringing it all together, *OECD Digital Education Outlook*, Paris: OECD Publishing, 2021.
13. K. Brennan and M. Resnick, New frameworks for 21st-century learning: Creating, collaborating, and computing, *IEEE Computer Graphics and Applications*, vol. 43, no. 2, pp. 22-29, 2023.
14. M. W. Loat, Agentic workflows in educational software: Beyond chatbots, in *Proc. IEEE Int. Systems Conf. (SysCon)*, 2024, pp. 12-18.
15. Y. Kim and H. Lee, Evaluation of student burnout and engagement using AI-based sentiment analysis, *Mobile Networks and Applications*, vol. 29, pp. 88-102, 2024.