

PERSONALIZED EDUCATION WITH SEMANTIC MEMORY & RAG

Presented by:

Basanth Periyapatna Rooper Kumar
Manav Rajesh Anandani
Nandhakumar Apparsamy
Nischitha Nagendran
Nitya Rondla

Project Advisors:

Dr. Simon Shim
Dr. Lee C. Chang

Introduction

The Personalized Education Agent addresses the "personalization gap" in e-learning by converting static academic materials into an adaptive, document-grounded AI tutor. Standard LMS platforms and general AI assistants often lack access to a learner's specific course materials, causing context-blind responses and hallucinated explanations. Our system combines Retrieval-Augmented Generation, ChromaDB-based semantic memory, LlamaParse document ingestion, and an adaptive Root-Cause Analysis engine. The final prototype generates personalized learning paths, supports grounded Socratic tutoring, visualizes mastery through a ReactFlow concept graph, and inserts remedial milestones when prerequisite gaps are detected. Final evaluation showed hallucinations reduced from 15.4% to <2%, a 0.94 RAGAS faithfulness score, 680ms average time-to-first-token, and 94% RCA diagnostic accuracy.

Methodology

The learner uploads course material and a learning goal, which initiates the pipeline. Documents are parsed using LlamaParse to preserve structure such as tables and formulas. The content is split using semantic chunking, and each chunk is converted into embeddings and stored in ChromaDB.

At query time, relevant chunks are retrieved using similarity search (Top-K) and passed into a RAG pipeline to generate grounded responses. The system uses low-latency LLM inference to ensure real-time interaction.

A Root-Cause Analysis (RCA) engine analyzes quiz performance to identify missing prerequisite concepts. Based on this, the system dynamically updates the learning path by inserting refresher milestones.

Outcome: Static PDFs are transformed into an adaptive, personalized, and document-grounded learning experience.

Analysis and Results

Introduction: Context Blindness in AI Tutors

Your Most e-learning platforms and general AI tutors do not understand a student's exact course materials, professor-specific notation, or current knowledge gaps. As a result, students often receive generic answers that may not match their uploaded textbooks, lecture notes, or class expectations.

This creates three core problems:

One-size-fits-all learning: static learning paths ignore individual mastery.


Hallucination risk: AI may generate unsupported or incorrect explanations.

Cold-start burden: students must manually turn dense PDFs into study plans.

Goal: Build an AI tutor that reads, remembers, retrieves, teaches, and adapts using the learner's own curriculum paragraph text

Technology Stack Section

The system uses a hybrid cloud-local architecture designed for grounded tutoring, low-latency interaction, and privacy-aware document retrieval.



Intelligent Features Section

Ingestion Hub: Upload materials + goal.
AI Tutor: Grounded lessons and Q&A.
Concept Map: Mastery, gaps, dependencies.
RCA Panel: Finds prerequisite gaps.
Linear Syllabus: Personalized milestones.
Quiz: Captures mastery signals.
Flow: Upload → Path → Learn → Quiz → Gap → Adapt.

Summary/Conclusions

The Personalized Education Agent shows that AI tutoring can move beyond generic Q&A into grounded, adaptive learning support. By combining RAG, semantic memory, LlamaParse ingestion, ReactFlow mastery visualization, and RCA-based remediation, the system reduces hallucinations while improving explainability. Key takeaway: A student's private curriculum can become an interactive learning environment instead of a static collection of PDFs.

Key References

- [1] Bloom, B. S. (1956). Taxonomy of educational objectives: The classification of educational goals. Handbook I: Cognitive domain. Longmans, Green.
- [2] Google Research. (2025). Learn Your Way: Reimagining textbooks with generative AI. Google Research Blog.
- [3] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474.
- [4] Li, Z., & Wang, X. (2025). Retrieval-augmented generation for educational application: A systematic survey. Computers and Education: Artificial Intelligence, 8, Article 100392.

Acknowledgements

We thank Dr. Simon Shim and Dr. Lee C. Chang for their guidance, technical feedback, and support throughout the DATA 298A/B project lifecycle. We also thank the Department of Applied Data Science at San José State University for providing the academic structure and resources needed to complete this capstone project.

Interested in partnering with us on a capstone project or exploring other ways to collaborate with the College of IDS? Contact CIDS_CAP@sjsu.edu to discuss potential projects and partnership opportunities.