# Personalized Education with Semantic Memory & RAG

**A full-stack AI tutoring platform that grounds responses in student-uploaded course materials, generates adaptive learning paths, and diagnoses prerequisite knowledge gaps.**


**Applied Data Science Department**

**Basanth Periyapatna Roopa Kumar, Manav Rajesh Anandani, Nandhakumar Apparsamy, Nischitha Nagendran, Nitya Rondla**

**Project Advisors: Dr. Simon Shim, Dr. Lee C. Chang**



## Abstract

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**.


## 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

Scan to share your learning challenges

## Technology Stack Section

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

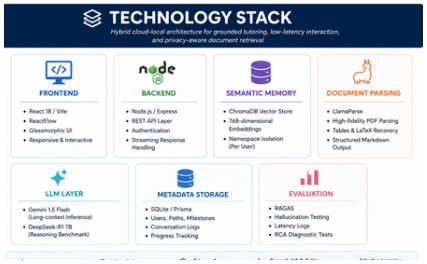


## Methodology

**Upload → Parse → Chunk → Embed → Retrieve → Generate → Adapt**

The learner uploads course material and a goal. The system parses documents with LlamaParse, stores semantic chunks in ChromaDB, retrieves relevant context through RAG, generates grounded tutoring responses with Gemini, and uses RCA to add refresher milestones when prerequisite gaps are detected.

Outcome: Static PDFs become an adaptive, document-grounded learning path.

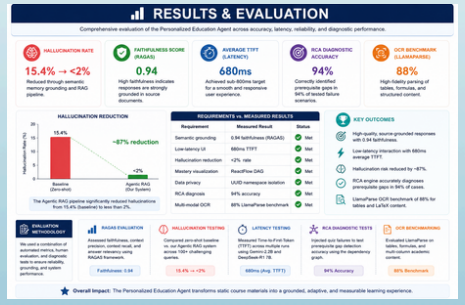


## Intelligent Features Section

**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.**

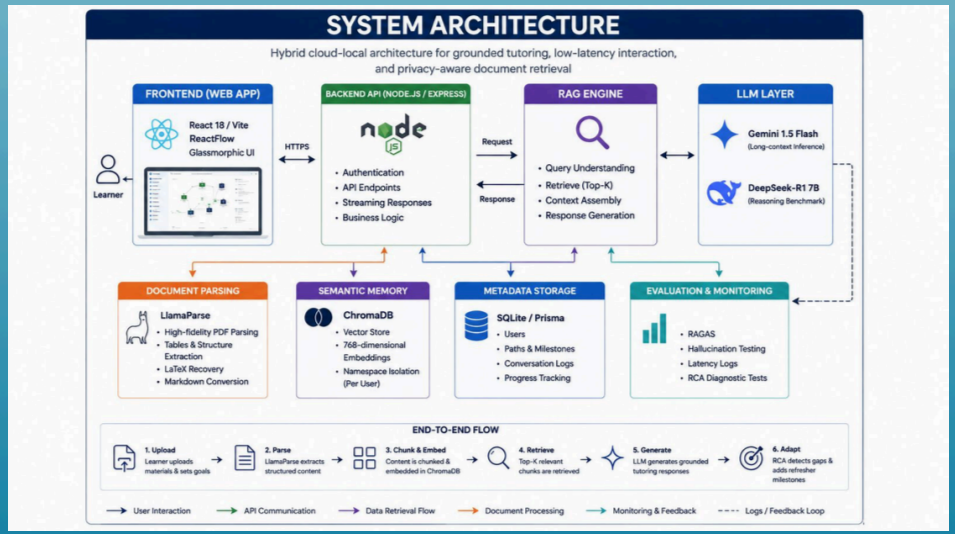


## Conclusion

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.

## Future Work

Future development should focus on:

**LMS Integration:** Connect with Canvas or Moodle through LTI 1.3.

**Multi-Agent Tutoring:** Add agents for math, coding, diagrams, and exams.

**On-Device Inference:** Improve privacy and reduce API dependency.

**Diagram-Aware Retrieval:** Link diagrams with surrounding text and explanations.

**Larger Evaluation:** Test learning gains with bigger classroom pilots.


## 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.


SCAN ME

Try yourself
Learn yourself