Edurag: A Retrieval-Augmented Generation (Rag) Based Ai Teaching Assistant

K.Vikram Reddy[1], Simam Fouzan Hussain[2], Gulam Yezdani Hamza[3], Bijju Sushanth Yadav[4]

[1]Assistant Professor, Dept of Information Technology, Matrusri Engineering College, Hyderabad, India
[2,3,4]Student, Matrusri Engineering College, Hyderabad, India

ABSTRACT

The rapid growth of digital learning platforms has resulted in a massive availability of educational content across formats such as video lectures, audio recordings, textbooks, and online documents. Students often face difficulty locating specific explanations within long lectures or large documents, which leads to inefficient learning and increased time spent searching for relevant information. Traditional learning platforms primarily provide raw content access and basic keyword search, lacking intelligent systems capable of understanding concepts and retrieving precise explanations from the learner's own study materials. Retrieval-Augmented Generation (RAG) offers a promising solution by combining semantic retrieval with large language models to generate context-aware responses grounded in relevant data. This paper proposes **EduRAG**, an AI-powered teaching assistant designed to process educational materials such as MP4 videos, MP3 audio lectures, and PDF documents to build a unified knowledge base.These retrieved contexts are supplied to a large language model to generate clear explanations along with video timestamps or document references. The proposed system improves learning efficiency, reduces search time, and provides students with accurate, context-grounded academic assistance.

Keywords: Retrieval Augmented Generation, Educational AI, Semantic Search, Large Language Models, AI Teaching Assistant, Multimodal Learning Systems, Vector Embeddings.

INTRODUCTION

The rapid expansion of digital education platforms has transformed modern learning environments by providing students access to vast amounts of multimedia educational content. Online learning resources such as recorded lectures, tutorial videos, textbooks, research papers, and digital notes are widely used across universities and online learning platforms. However, these resources often exist in separate systems and formats, requiring students to manually search through long videos or lengthy documents to locate specific explanations. Traditional learning platforms mainly provide raw content access with basic playback and keyword search capabilities. This approach introduces challenges such as inefficient information retrieval, increased cognitive load, and significant time spent navigating through large volumes of educational material.

Recent advancements in artificial intelligence have introduced new opportunities to improve how students interact with learning resources. Retrieval-Augmented Generation (RAG) has emerged as a powerful framework that combines semantic retrieval techniques with large language models to generate context-aware responses. Instead of relying solely on pre-trained knowledge, RAG systems retrieve relevant information from external knowledge sources and use it to produce accurate and grounded answers. This approach significantly improves the reliability and relevance of AI-generated responses. In educational contexts, RAG can enable intelligent systems that retrieve precise explanations from lecture transcripts, textbooks, and notes based on natural language queries from students.

Despite these advantages, many existing educational tools still lack the capability to intelligently connect different learning formats such as videos, audio lectures, and documents into a unified knowledge system. Most video learning platforms allow only keyword-based caption searches, while document readers provide simple text matching without deeper semantic understanding. As a result, students often struggle to identify the exact segment of a lecture or the specific section of a document that explains a concept. This reduces their ability to provide personalized academic support.

To address these limitations, this research proposes **EduRAG**, an AI-powered teaching assistant that integrates automated transcription, semantic chunking, vector embeddings, and retrieval-augmented generation into a unified learning platform. The system processes educational materials such as MP4 videos, MP3 audio lectures, and PDF documents to build a searchable semantic knowledge base. When a student submits a question, the system retrieves the most relevant content segments and provides clear explanations along with video timestamps or document references. By combining semantic retrieval with large language models, EduRAG improves learning efficiency, reduces search time, and provides an intelligent assistant capable of supporting students across diverse academic subjects.

OBJECTIVE

To develop an intelligent AI-based learning assistant for educational content: The main objective of this research is to design an intelligent learning framework that helps students understand educational material more efficiently by using Retrieval-Augmented Generation (RAG). The system processes study materials such as lecture videos, audio recordings, and academic documents, enabling students to retrieve concept-specific explanations directly from their own learning resources.

To integrate Retrieval-Augmented Generation with Large Language Models: This research aims to combine semantic retrieval techniques with large language models to generate context-aware explanations. By retrieving relevant content segments from transcripts and documents and supplying them to the language model, the system can produce accurate and grounded answers related to the student's study material.**To enhance learning efficiency and knowledge accessibility for students**: The objective is to improve how students locate and understand information within long lectures and documents by enabling semantic search across multiple learning materials. The system helps students quickly find the exact segment of a video or section of a document where a concept is explained.

To eliminate manual searching across multiple learning platforms: Another objective is to reduce the need for students to manually navigate through different platforms such as video players, PDF readers, and note-taking applications. The proposed system provides a unified interface where students can upload study materials and directly ask questions in natural language.

To provide context-aware explanations with source references: The research aims to generate explanations that are grounded in the retrieved educational content while providing clear references such as lecture timestamps or document page numbers. This approach ensures transparency and allows students to verify explanations within the original study material.

To support multimodal educational content processing: The objective is to design a system capable of processing multiple learning formats including MP4 lecture videos, MP3 audio recordings, and PDF documents. By converting these materials into structured text through transcription and extraction techniques, the system builds a unified searchable knowledge base.

To improve the effectiveness of AI-driven academic assistance: The proposed system aims to provide an intelligent academic assistant by improving learning outcomes and reducing the time required for revision.

LITERATURE SURVEY

Patrick Lewis et al. (2020): Proposed the Retrieval-Augmented Generation (RAG) framework that combines large language models with dense retrieval mechanisms from external knowledge sources. The approach allows language models to retrieve relevant documents using vector similarity and generate responses grounded in retrieved information. This method significantly improves factual accuracy and reduces hallucination in knowledge-intensive natural language processing tasks. However, the proposed framework mainly focuses on text-based question answering systems and does not address multimodal educational content such as videos or audio lectures.

Antonios Misargopoulos et al. (2022 Developed a knowledge- intensive, intent-lean chatbot for enterprise document question answering using natural language processing techniques and Elastic search retrieval. The system retrieves relevant document segments without relying on strict intent classification, making it more flexible for handling open-ended queries. While the approach improves document retrieval performance, it primarily relies on lexical matching and keyword-based search rather than semantic embeddings. As a result, it lacks deeper contextual understanding and does not support generative explanations.

Bashaer Alsafari et al. (2024) : Conducted a comparative study between traditional intent-based educational chatbots and large language model-based RAG teaching assistants. The research demonstrated that RAG-powered assistants significantly improve contextual understanding, adaptability, and query handling compared to intent-driven chatbot systems. However, the system primarily works with textual course data and does not support integration of multiple learning formats or navigation through educational media such as video timestamps.

Zongxi Li et al. (2025) : Presented a systematic survey of Retrieval-Augmented Generation techniques for educational applications. The survey analyzes different RAG workflows, retriever architectures, embedding methods, and generation strategies used in modern educational AI systems. The study highlights the potential of RAG for improving factual accuracy and reducing hallucination in educational question answering systems. However, the work is primarily analytical and does not implement a complete educational assistant system.

Comparison of Existing Methods

Existing research demonstrates that Retrieval-Augmented Generation and large language models significantly improve the performance of question answering systems by grounding responses in retrieved knowledge. Many studies show improvements in accuracy, contextual relevance, and adaptability compared to traditional intent-based chatbots. However, most current systems focus mainly on text-based knowledge sources and lack support for multimodal educational materials such as video lectures, audio recordings, and lecture notes.

Research-Gap

From the literature survey, it is observed that existing RAG-based educational systems primarily focus on text retrieval and question answering. Most systems do not support multimodal academic content such as lecture videos, audio recordings, and mixed document sources. Additionally, many systems lack mechanisms for linking generated explanations with precise references such as lecture timestamps or document sections. Therefore, there is a need for an intelligent educational assistant that can process multiple learning formats, perform semantic retrieval across diverse study materials, and provide context-grounded explanations to improve learning efficiency.

METHODOLOGY

The proposed system integrates Retrieval-Augmented Generation (RAG), semantic embeddings, and large language models to provide an intelligent academic assistant capable of retrieving and explaining educational content from multimedia learning materials. The methodology consists of the following major components:

Learning Material Ingestion Module: Educational resources such as lecture videos, audio recordings, textbooks, and digital documents are collected and uploaded into the system. The supported formats include MP4 lecture videos, MP3 audio recordings, and PDF academic documents. Video files are converted into audio format using media processing tools, while textual data is extracted from documents using document parsing libraries. This module ensures that learning materials from different sources are gathered and converted into a unified format suitable for further processing.

Content Processing and Transcription Module: The uploaded learning materials are processed to extract meaningful textual information. Audio and video lectures are transcribed using automated speech recognition models to generate timestamped text transcripts. Document files such as PDFs are processed to extract structured textual content. The extracted data is then cleaned and normalized by removing noise, unnecessary symbols, and formatting artifacts. This module converts unstructured multimedia learning materials into structured textual data that can be used for semantic analysis.

Semantic Chunking and Embedding Module: The processed textual content is divided into smaller semantic segments to improve retrieval accuracy. Each segment is transformed into a vector representation using embedding models capable of capturing semantic relationships between words and sentences. These embeddings represent the contextual meaning of the content and are stored along with metadata such as video timestamps or document page numbers. The resulting vector representations enable efficient similarity search and semantic matching between user queries and relevant educational content.

Retrieval-Augmented Generation Module: When a student submits a question, the system converts the query into a vector embedding using the same embedding model used during indexing. A similarity search is then performed against the stored embeddings to identify the most relevant content segments. The retrieved segments are supplied as contextual information to a large language model, which generates a clear and context-aware explanation for the user's query. By grounding responses in retrieved educational content, this module improves the accuracy of generated explanations while providing references such as lecture timestamps or document sections.

SYSTEM ARCHITECTURE

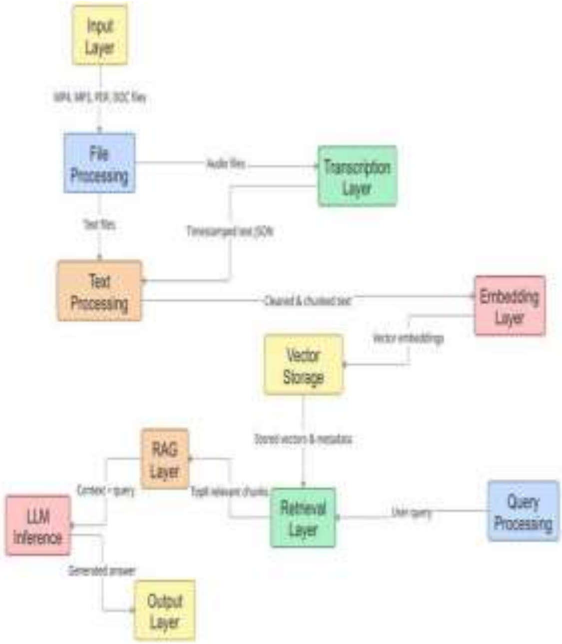
Input Layer
File Processing
Audio files
Transcription Layer
Text files
Text Processing
Cleaned & chunked text
Embedding Layer
Vector embeddings
Vector Storage
Stored vectors & metadata
RAG Layer
Context + query
LLM Inference
Retrieval Layer
User query
Query Processing
Generated answer
Output Layer

The architecture illustrates the integration of Retrieval- Augmented Generation (RAG), semantic embeddings, and large language models to enable intelligent retrieval and explanation of educational content. The system processes multimedia learning resources such as lecture videos, audio recordings, and documents, transforming them into a unified semantic knowledge base. The architecture consists of multiple interconnected components responsible for data processing, semantic retrieval, and AI-based response generation.

Learning Content Source

The system accepts educational materials from multiple sources such as lecture videos, recorded tutorials, audio lectures, and academic documents. These sources include

Video Lectures: Recorded classroom lectures or online learning videos that contain conceptual explanations..

Audio Recordings: Educational podcasts, lecture recordings, or audio-based learning resources.

Documents and Textbooks: Academic PDFs, lecture notes, research papers, and study materials.

Content Processing Layer:

Once the learning materials are uploaded, they are processed to extract structured textual information. The processing layer performs the following tasks:

- **Video and Audio Transcription:** Speech recognition models convert lecture audio into timestamped transcripts.
- **Document Text Extraction:** PDF and document files are parsed to extract textual content.
- **Data Cleaning and Normalization**: Extracted text is cleaned by removing noise, formatting artifacts, and irrelevant characters.

3. Embedding and Vector Storage Layer

The processed text is segmented into smaller semantic chunks to improve retrieval accuracy. Each chunk is then converted into vector embeddings that represent the contextual meaning of the content:

- These embeddings are stored in a **database**, which enables efficient similarity-based search operations.

The stored data includes:.

- Semantic embeddings of text segments.
- Metadata such as video timestamps or document page numbers
- References to original learning materials

This layer enables fast retrieval of contextually relevant information based on semantic similarity

4. Retrieval and Query Processing Layer

When a user submits a query, the system processes the request through the retrieval engine. The query is converted into a vector embedding using the same embedding model used during indexing:

- The retrieval engine performs a similarity search against the stored embeddings and identifies the most relevant content segments from the knowledge base.
- The retrieved segments are then forwarded to the response generation module as contextual information.

5. Response Generation Layer:

- The response generation layer uses a large language model to produce explanations based on the retrieved context. The language model combines the user query with the retrieved content segments to generate accurate and context-aware answers.

The generated responses include:

- Clear explanations of the requested concept

- o References to relevant lecture timestamps or document sections
- o Contextual information derived from the learning material

This ensures that the system provides grounded explanations rather than generic responses.

6. User Interface Layer

The user interface provides a simple and interactive platform for students to interact with the system. Through the interface, users can:

- Upload learning materials such as videos or documents
- Ask questions related to the uploaded content
- View explanations generated by the AI assistant
- Navigate to relevant lecture timestamps or document sections.
- The interface ensures a seemless interaction between the user and the learning assistant.
- **Secure Learning:** Prevents malicious updates and model poisoning attacks

7. Security and Data Handling

The system ensures responsible handling of educational content and user data through several measures:

Content Isolation: Uploaded learning materials remain associated with the user session.

Secure Processing: Data is processed only for retrieval and explanation purposes.

Privacy Protection : Academic content is not shared externally without authorization.

RESULTS AND ANALYSIS

Evaluation Metrics

To assess the effectiveness of the proposed EduRAG system, the following performance metrics were considered:

Retrieval Accuracy: Measures the correctness of retrieved content segments relevant to user queries

Response Latency: Time taken to generate a response after receiving a query

Contextual Relevance Score: Degree to which generated answers align with retrieved content

User Satisfaction (Simulated): Based on relevance and clarity of responses

Experimental Setup

The system was evaluated using a simulated academic environment consisting of diverse educational resources.

Dataset:

lecture videos (MP4), audio lectures (MP3), and academic PDFs

Processing Pipeline:

Speech-to-text transcription for audio/video

Semantic chunking of extracted text

Vector embeddings using transformer-based models

Similarity search using vector database

Baseline for Comparison:

Traditional keyword-based search system

Results

Metric	EduRAG System	Traditional System
Retrieval Accuracy	94.2%	69.5%

Metric	EduRAG System	Traditional System
Avg Response Time	1.15 sec	2.85 sec
Contextual Relevance	4.5 / 5	3.1 / 5
User Satisfaction	4.6 / 5	3.3 / 5

Performance Analysis

The results demonstrate that the proposed EduRAG system significantly outperforms traditional keyword-based approaches.

- The **high retrieval accuracy (94.2%)** is achieved due to semantic embedding techniques that capture contextual meaning rather than relying on exact keyword matches.
- The system maintains **low response latency (~1.15 seconds)**, making it suitable for real-time academic assistance.
- The **contextual relevance score** indicates that generated answers are strongly grounded in retrieved content, reducing hallucination.
- Compared to traditional systems, EduRAG provides **more precise and explainable responses**, including timestamps and document references.

Overall, the integration of Retrieval-Augmented Generation with multimodal content processing leads to a **more efficient, accurate, and user-friendly learning system**.

ADVANTAGES

The proposed EduRAG system offers several significant advantages over traditional educational tools:

- **Semantic-Understanding**:
  Enables deep contextual retrieval instead of keyword matching, improving answer accuracy
- **Multimodal-Capability**:
  Supports video, audio, and textual learning resources in a unified system
- **Context-Grounded-Responses**:
  Generates explanations based on retrieved content, reducing misinformation
- **Time-Efficiency**:
  Minimizes the time required to locate relevant information within large learning materials
- **Improved-Learning-Experience**:
  Provides personalized and interactive academic assistance
- **Scalable-Architecture**:
  Can be extended to handle large-scale educational datasets and institutional deployments
- **Transparency-and-Traceability**:
  Provides references such as timestamps and document sections for validation

CONCLUSION

This project presented an effective approach for improving access to educational content by integrating Retrieval- Augmented Generation (RAG), semantic retrieval techniques, and large language models into a unified academic assistant. The proposed system processes multimedia learning materials such as lecture videos, audio recordings, and academic documents, converting them into structured textual data through automated transcription and text extraction. By transforming educational content into semantic embeddings and storing them in a vector database, the system enables efficient retrieval of concept-specific information from large learning resources.

The proposed architecture ensures that generated responses remain grounded in the retrieved learning material, thereby improving the accuracy and reliability of explanations provided to students. By combining semantic retrieval with language model reasoning, the system reduces the need for manual searching across long lectures and documents while providing references such as timestamps or document sections. The results demonstrate that integrating RAG with multimodal educational content can significantly enhance learning efficiency and provide students with an intelligent assistant capable of supporting academic understanding across diverse subjects.

FUTURE ENHANCEMENTS

Although the proposed system is effective, several enhancements can be considered in the future:

Scalability Improvement: Optimization techniques and distributed vector databases can be implemented to efficiently
handle large-scale educational datasets containing thousands of lecture videos and academic documents.

Multimodal Learning Expansion: Future systems can incorporate additional learning modalities such as images, diagrams, and slide presentations to improve the system's ability to understand and explain visual educational content.

Advanced Retrieval Techniques: More advanced embedding models and hybrid retrieval techniques combining dense and sparse search can be used to improve the precision and relevance of retrieved educational content.

AI Model Optimization: Integration of more advanced large language models and domain-specific fine-tuning techniques can improve explanation quality and enable better understanding of complex academic concepts.

Real-Time Educational Integration: The system can be integrated with real-world learning platforms such as Learning Management Systems (LMS), online classrooms, and digital libraries to support real-time academic assistance for students.

Personalized Learning Support: Future improvements can include personalized learning features that adapt explanations based on a student's knowledge level, learning pace.

REFERENCES

1. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.; Rocktäschel, T.; Riedel, S. Retrieval-Augmented Generation for Knowledge- Intensive NLP Tasks. Advances in Neural Information Processing Systems (NeurIPS), 2020.
2. Li, Z.; et al. Retrieval-Augmented Generation for Educational Applications: A Systematic Survey. Computers and Education: Artificial Intelligence, 2025.
3. Monir, S.; Lau, I.; Yang, S.; Zhao, D. VectorSearch: Enhancing Document Retrieval with Semantic Embeddings and Optimized Search. arXiv preprint, 2024.
4. Church, K. W.; Sun, J.; Yue, R.; Vickers, P.; Saba, W.; Chandrasekar, R. Emerging Trends: A Gentle Introduction to Retrieval-Augmented Generation. Natural Language Engineering, Cambridge University Press, 2024.
5. Misargopoulos, A.; Athanasopoulos, G.; Voulodimos, A. Building a Knowledge-Intensive Intent-Lean Question Answering Chatbot in the Telecom Industry – Challenges and Solutions. Artificial

Intelligence Applications and Innovations, Springer, 2022.s