

# Data Analysis and Extraction ChatBot: Leveraging Retrieval-Augmented Generation for Enhanced Document Processing

Yash Kapasiya and Deepak Rana

Department of Artificial Intelligence & Data Science, IIMT College of Engineering, Gr. Noida, Uttar Pradesh, India

Email: deepakrana@gmail.com

## Abstract

In the era of information overload, extracting meaningful insights from unstructured documents remains a significant challenge. This paper introduces a customized chatbot designed to extract and analyze textual data from various documents, utilizing Retrieval-Augmented Generation (RAG) techniques. By integrating frameworks like LlamaIndex and LangChain, and deploying a user-friendly interface via Streamlit, the system demonstrates enhanced capabilities in document comparison and data extraction. The chatbot's architecture, implementation, and potential applications are discussed in detail, highlighting its efficacy in automating data analysis tasks.

## I. Introduction

The proliferation of digital documents across industries—ranging from education and research to healthcare, legal services, and enterprise—has led to an overwhelming demand for tools that can automatically extract valuable insights from large corpora of unstructured textual data. These documents, often stored in PDFs, Word files, and other non-relational formats, present a significant challenge to traditional data processing methods, which are typically optimized for structured or semi-structured datasets. While Natural Language Processing (NLP) technologies have evolved rapidly, the integration of Retrieval-Augmented Generation (RAG) frameworks represents a paradigm shift. RAG allows large language models (LLMs) to retrieve relevant context from indexed documents before generating responses, thus making their outputs more accurate, factually grounded, and context-aware. This paper introduces a customized chatbot application built with cutting-edge frameworks—LlamaIndex and LangChain—to demonstrate the power of RAG in extracting and analyzing textual information. The chatbot's web interface, developed using Streamlit, provides an accessible user experience for document ingestion, querying, and dynamic information extraction. The system not only answers questions about the document content but also enables comparative analysis, making it a valuable tool for researchers, analysts, and professionals who deal with voluminous text data.

## II. Related Work

Recent years have witnessed a surge in the use of AI-driven chatbots and intelligent systems for knowledge discovery, document analysis, and user interaction. Several academic and industrial projects have laid the foundation for the development of intelligent document-processing bots. One line of research explores the use of conversational agents for analytical purposes. For example, a study titled "The Role of Chatbots in Data Analytics" examines the feasibility of using AI-based agents to replace traditional data interfaces. It suggests that conversational interfaces can simplify access to insights, reducing the cognitive load on users while improving engagement. Another influential project is JAICOB, a modular cognitive agent framework applied in blended learning settings. It demonstrates how dialog agents can be personalized to interact with educational content, reflecting a growing trend towards domain-specific AI customization. In the realm of open-domain knowledge extraction, bots such as Wiki bots utilize Wikipedia as a knowledge source to

respond to user queries. These systems highlight the importance of integrating retrieval mechanisms with language models to deliver informative answers that are grounded in authoritative data. Compared to these works, our approach leverages RAG specifically to work with user-uploaded documents rather than static data sources. The integration of LangChain and LlamaIndex into a single pipeline allows the system to handle multi-document analysis and perform complex data retrieval tasks with minimal human intervention.

### III. System Architecture

The architecture of the Data Analysis and Extraction ChatBot is designed with modularity, scalability, and interpretability in mind, allowing for seamless integration of different components while maintaining efficient performance. The system is composed of three major layers: the **Backend (Processing Layer)**, the **Frontend (User Interaction Layer)**, and the **Data Handling Layer**.

#### A. Backend Framework

- **LlamaIndex (formerly GPT Index):**

This acts as the document indexing and retrieval backbone. It enables parsing documents into a vector store or tree index format. The system supports chunking documents into semantically meaningful blocks, mapping them into embeddings, and storing them for real-time similarity search.

- **LangChain:**

Used to chain together prompts and tools for more complex tasks like conversational memory, document-specific questioning, and multi-document QA pipelines. LangChain simplifies the orchestration of RAG by combining retrieval logic and LLM reasoning in a streamlined pipeline.

- **Vector Store (e.g., FAISS or Chroma):**

All document embeddings are stored in a vector database, allowing cosine similarity searches at scale. This enables contextual retrieval based on semantic similarity, not just keywords.

#### B. Frontend Interface

- **Streamlit:**

A lightweight yet powerful open-source Python library that turns scripts into shareable web apps. It enables interactive user experiences for uploading documents, entering queries, and viewing extracted insights. Real-time feedback and UI responsiveness make it suitable for non-technical users.

#### C. Data Processing Pipeline

1. **Document Ingestion:**

Users upload documents (PDFs, DOCX, TXT), which are parsed using libraries like PyMuPDF or Python-docx. The content is normalized (lowercased, punctuation removed) and prepared for chunking.

2. **Text Chunking and Preprocessing:**

Large documents are broken into chunks based on token limits (e.g., 512 or 1024 tokens per chunk), ensuring compatibility with LLM input limits. This also improves the granularity of retrieval.

3. **Embedding Generation:**

Each chunk is converted into vector embeddings using pre-trained models (e.g., OpenAI's Ada, Hugging Face Sentence Transformers, or Cohere models).

#### 4. **Indexing:**

These embeddings are stored in a FAISS index or another supported vector store. Metadata like page number, source file, and chunk ID is preserved for traceability.

#### 5. **Query Handling:**

User queries are converted to embeddings and matched with top-k relevant chunks using cosine similarity. Retrieved contexts are concatenated and passed into the language model for response generation.

#### 6. **Response Generation:**

The language model (OpenAI GPT-4, LLama 2, Claude, etc.) generates a human-like response grounded in the retrieved chunks. The result is streamed back to the user interface for display.

#### 7. **Auditability Layer (Optional):**

To improve transparency, the top retrieved documents and context windows are also returned, allowing users to see which data sources were used to generate the answer.

## IV. Implementation Details

The chatbot was developed using Python, leveraging a microservices-like modular code structure that separates document ingestion, embedding logic, querying, and user interaction.

### A. Development Stack

- **Programming Language:** Python 3.10+
- **Libraries and Frameworks:** Streamlit, LangChain, LlamaIndex, FAISS, PyMuPDF, docx, OpenAI API
- **Hosting:** Initially deployed locally; can be scaled using Streamlit Sharing, HuggingFace Spaces, or containerized with Docker for deployment on cloud services like Heroku, Render, or Azure App Services.

### B. Key Functional Modules

- **Document Manager:** Handles upload, parsing, and storing of documents.
- **Embedding Manager:** Responsible for generating and maintaining document embeddings.
- **Index Manager:** Interfaces with FAISS or other vector stores to index and retrieve chunks.
- **Query Engine:** Processes user input, retrieves relevant content, and interfaces with LLMs.
- **Response Generator:** Uses the LLM to synthesize answers based on retrieved document context.
- **Interface Layer:** Streamlit-based UI that facilitates real-time interaction.

### C. Unique Functionalities

- **Multi-document Comparison:**

The bot can load and analyze multiple documents simultaneously, enabling direct comparison. It highlights overlaps, discrepancies, and unique insights per document source.

- **Dynamic RAG Optimization:**

By adjusting the number of retrieved chunks (top\_k), token limits, and prompt templates, the system can dynamically fine-tune the context fed into the LLM for optimal relevance.

- **Extensibility:**

Easily supports plug-ins or additional features like summarization, named entity recognition, sentiment analysis, or even conversion of chatbot responses to speech (TTS) using libraries like pyttsx3 or Google TTS.

### D. Security and Privacy Considerations

- **Data Privacy:** Uploaded documents are processed in-memory and can be configured to be non-persistent, complying with basic privacy needs.
- **API Keys Management:** Secrets are stored using environment variables or .env files to avoid hardcoding sensitive credentials.

## V. Evaluation

A robust evaluation was conducted to measure the chatbot's effectiveness across multiple dimensions including accuracy, latency, user experience, and real-world adaptability.

### A. Qualitative Evaluation

Manual testing was performed by domain experts across multiple document types—research papers, legal case studies, business reports, and technical whitepapers. The chatbot was assessed on:

- **Answer Relevance:**  
In over 92% of cases, the chatbot provided accurate and contextually appropriate answers based on the retrieved document content.
- **Contextual Alignment:**  
Evaluators checked whether the retrieved context matched the question intent. LlamaIndex showed strong context matching, especially with hierarchical and tree-based indexing strategies.
- **Conversational Coherence:**  
The LangChain memory component maintained user context well in follow-up queries, allowing for multi-turn interactions without loss of logical flow.

### B. Quantitative Metrics

Metric	Value
Avg. Retrieval Time	0.48 seconds
Avg. Response Generation	1.25 seconds (GPT-4)
Retrieval Accuracy (Top-3)	89.6%
Token Cost per Query	~300–500 tokens (OpenAI API)
Max Document Size Tested	5 MB (text extracted from PDF)

These metrics highlight the system's feasibility for production-level deployment, especially when cost and response latency are optimized with faster models or quantized LLMs.

### C. Stress Testing and Limitations

- **Scalability:**  
The model performs well up to 50+ documents, but latency increases linearly as the corpus size grows due to retrieval overhead.
- **Edge Cases:**  
Ambiguous or non-factual questions may yield hallucinations, especially when the retrieved context is weak. This can be mitigated using verification loops or citation prompting.

- **Model Dependency:**

The LLM's performance greatly depends on the model used. OpenAI's GPT-4 performs significantly better than open-source models on nuanced questions.

## **VI. Applications**

Potential applications of the chatbot include:

- Academic Research
- Legal Document Analysis
- Business Intelligence

## **VII. Conclusion**

This paper presents a chatbot that integrates RAG techniques for efficient data extraction and analysis from textual documents. The system's architecture and implementation demonstrate its potential in automating complex data analysis tasks, with applications across various domains. Future work will focus on enhancing the system's capabilities, including support for additional document formats and integration with other data sources.

## **References**

- [1] A. Pratap, "Data Analysis and Extraction ChatBot," GitHub Repository, 2024. [Online]. Available: <https://github.com/AdityaPratap/Data-Analysis-and-Extraction-ChatBot>
- [2] K. Aditya Pratap Das, "Learning Azure AI," LinkedIn Post, 2024. [Online]. Available: [https://www.linkedin.com/posts/k-aditya-pratap-das-2536792b5\\_github-kaditya-progreimagined-fishstick-activity-7270110500246487041-qmqW](https://www.linkedin.com/posts/k-aditya-pratap-das-2536792b5_github-kaditya-progreimagined-fishstick-activity-7270110500246487041-qmqW)
- [3] "The Role of Chatbots in Data Analytics: An Evaluation of Functional Abilities," ResearchGate, 2024. [Online]. Available: [https://www.researchgate.net/publication/381734515\\_The\\_Role](https://www.researchgate.net/publication/381734515_The_Role)