International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** VI **Month of publication:** June 2025

DOI: https://doi.org/10.22214/ijraset.2025.72560

"Enhancing Equity Research with AI: A LangChain-Based News Analysis Framework"

Thejaswini K S[1], Yashwanth N[2], Vivek D[3], Chandana H S[4], Mrs. Sindhu K S[5]
[1, 2, 3, 4]*Department of Information Science Engineering (Malnad College of Engineering, Hassan), Karnataka, India*
[5]*Assistant Professor of Malnad College of Engineering, Hassan*

Abstract: The advancement in chatbot technology and large-scale data processing has significantly transformed financial analysis and equity research and many other things. This research presents a LangChain-based framework to process financial documents, including company reports and market trends, to generate actionable investment insights. By integrating state-of-the-art technologies such as Large Language Models (LLMs) and advanced Natural Language Processing (NLP) techniques, the framework will support efficient data extraction, semantic analysis, and response generation. The proposed system begins by extracting key information from PDF documents, such as company financials and stock performance reports, which are then processed and segmented into manageable text chunks. These chunks are embedded into high dimensional vectors using techniques like Word2Vec or Doc2Vec, allowing the system to capture semantic relationships and store them in a semantic knowledge base. The knowledge base is further enhanced with tools like FAISS for efficient similarity search and information retrieval. In the second phase, the system responds to user queries by analyzing the context and questions posed. It retrieves relevant information from the knowledge base and generates responses using a generative AI model like GPT, ensuring high relevance and accuracy. The research also compares the system's efficiency in answering various types of investment-related questions, showcasing the chatbot's capability to assist users in making informed decisions. The framework's versatility and scalability, supported by cutting edge AI models and semantic search, demonstrate its potential to revolutionize the way equity research is conducted, providing financial analysts and investors with a more efficient and accessible tool for market analysis in the digital age.
Keywords: Equity Research, Natural Language Processing, Artificial Intelligence, Chatbot, Lang chain, Generative AI, LLM, IR, TF-IDF, FAISS, Similarity Search.

I. INTRODUCTION

In the digital age, the overwhelming volume of news content available online has made it increasingly difficult for individuals to effectively process and understand information. This has created a demand for advanced technologies capable of extracting valuable insights from vast amounts of news data. The proposed project introduces a LangChain-powered, Large Language Model (LLM)-based news research tool that allows users to interact with news articles by generating concise summaries and answering context specific queries based on the article's content.

The project draws inspiration from equity research chatbots that leverage LLMs to analyze financial documents and generate insights. Similarly, this system utilizes the LangChain framework to process news content, employing modular text loaders, recursive chunking, and vector databases for efficient document handling. These components enable the system to encode and transform text into searchable vectors, allowing users to retrieve relevant information based on specific queries. By using NLP and generative AI techniques, the system ensures accurate, context-sensitive, and language-specific answers, making it a powerful tool for enhancing news analysis.

This research highlights the potential of combining LLMs with advanced information retrieval techniques to improve the efficiency and accuracy of news content analysis. The proposed tool offers an intuitive user interface that simplifies the interaction between users and complex news data, providing both summaries and detailed answers to enhance user understanding of news topics.

II. LITERATURE REVIEW

Recent advancements in Natural Language Processing (NLP) and Large Language Models (LLMs) have paved the way for innovative applications in equity research. Several studies have explored the use of NLP in analyzing financial data and generating investment recommendations. For instance, research on stock prediction using NLP has shown that LLMs can effectively analyze financial texts like equity reports and news articles, offering personalized and accurate recommendations for investors [1][2].

Studies on equity research chatbots emphasize the need for real-time adaptability and personalized advice. While previous models used LSTM and RNN to predict stock movements, they lacked real-time customization based on individual user needs [3][4]. In contrast, your project aims to bridge this gap by creating a chatbot that not only provides accurate stock recommendations but also adapts to specific user queries in real-time. Furthermore, research on generative AI and LLMs highlights their ability to process vast amounts of financial data, providing context-sensitive responses. These technologies are essential for developing a chatbot that can analyze news articles and provide relevant, up-to-date insights for investors [5][6]. Studies on user engagement in financial chatbots stress the importance of intuitive interfaces and user trust. By integrating these aspects, your project aims to enhance user experience and ensure effective financial decision-making through an interactive, real-time investment tool [7][8].

III. METHODOLOGY

The core objective of this methodology is to create an AI-driven chatbot that can assist in equity research by interpreting financial documents and providing contextual responses. The proposed approach is structured into two main phases. The first phase focuses on building a robust Knowledge Base by processing raw financial data, and the second phase deals with interpreting user queries to generate meaningful responses. The architecture guiding this process is depicted in Fig. 2.

1) Phase I: Knowledge Base Creation

Step 1 – Data Extraction from PDF Documents: Initially, financial and equity research data is ingested into the system in the form of PDF files. These documents may contain company performance summaries, earnings reports, or market analysis. To retrieve the text content, we utilize the PyPDF Python library, which allows us to parse through and extract readable data from each page of the file. The text is prepared and cleaned to serve as the raw material for downstream processing.

Step 2 - Document Chunking with the help of vector: Following extraction, the large text corpus is segmented into smaller blocks. This is necessary because our LLM, Gemini Pro 1.0, has a token limitation of 32,760. To maintain processing efficiency and accuracy, the text is divided into chunks of approximately 10,000 characters. This chunking ensures compatibility with the model's input constraints and enables targeted processing of each segment.

Step 3 – Semantic Embedding of Chunks: Each text chunk is transformed into a numerical vector using embedding techniques such as Word2Vec or Doc2Vec. These embeddings represent the semantic relationships within the content and are essential for identifying key insights from financial documents. They capture meanings associated with financial performance metrics, market movements, and analyst commentary in a form that the model can understand and compare.

Step 4 – Indexing via Semantic Search (FAISS): The resulting embeddings are indexed using the FAISS library, which is optimized for fast similarity searches. FAISS organizes the vector space to allow efficient retrieval of related content by comparing distances between vectors. This clustering ensures that semantically similar pieces of information, such as quarterly reports or analyst opinions, are grouped together, making them easier to access during query processing.

Step 5 – Creation of the Knowledge Base: All the indexed data is compiled into a central repository called the Knowledge Base. This semantically structured database serves as the core reference point during user interaction. When a user inputs a query, the system scans this Knowledge Base to retrieve the most relevant information based on semantic proximity.

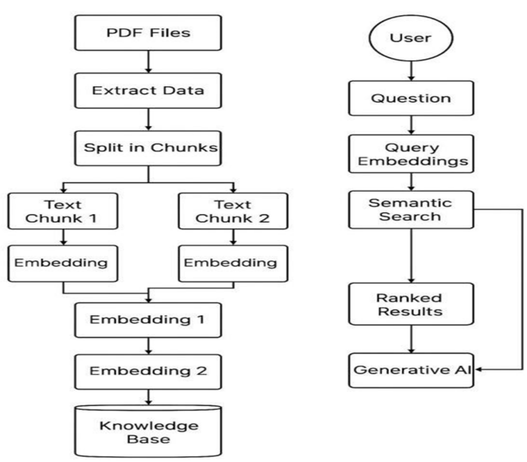

Fig. 1. Architecture Flow of Chatbot Development.

2) Phase II: User Query Response Generation

Step 1 – Prompt Intake and Segmentation: When a user interacts with the system, their input is interpreted as a prompt comprising two components: context and query. The context provides necessary background (e.g., financial ratios, historical trends), while the query specifies the user's question. This segmentation enhances the chatbot's ability to understand the full scope of the request.

Step 2 – Query Embedding and Chunking: The user query is processed similarly to the PDF data—it is converted into vector format. The vectorization process ensures that the user's question exists in the same semantic space as the stored knowledge, enabling an effective comparison. To optimize this comparison, the prompt is also broken into manageable chunks.

Step 3 – Similarity Matching and Retrieval: The embedded query is compared with the stored embeddings in the Knowledge Base using FAISS. This similarity search identifies the closest matches based on vector proximity. The system retrieves segments that are most semantically relevant to the user's prompt, ensuring contextual accuracy.

Step 4 – Generative AI Response Creation: After retrieving relevant content, a generative language model (such as GPT or Gemini Pro) is employed to construct a response. The model synthesizes the retrieved information along with the user's original query to deliver a coherent and domain specific answer.

Step 5 – Ranking and Refinement of Responses: Multiple response options may be generated, which are then ranked by their relevance and confidence scores. The top-ranked answer is presented to the user. Over time, the system uses feedback and interaction data to refine its search and generation mechanisms, continuously improving its accuracy and responsiveness.

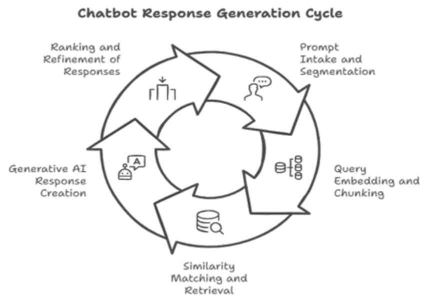

Fig. 2. Equity ChatBot lifecycle

IV. RESULTS AND DISCUSSIONS

By analyzing the text data extracted from PDF transcripts, we categorized the types of questions posed to the model into three groups:

1) Fact / Number-Based Questions: Answers to these questions are guaranteed to be correct when the information is present in the input files and is fact-based. For example, questions about the company's operating margin are answered accurately and succinctly.
2) Behavioral Brief Questions: These questions are only answered if management addressed them during the conference call. The responses are available in the PDF transcript. If the information is missing, the model explicitly states that management has not answered the question. For instance, questions about the company's plans to outperform competitors are only answered if such plans were discussed in the transcript.
3) Investment Advice Questions: If a user requests investment advice, the model refuses to respond, as providing such advice publicly is not appropriate according to government regulations. The model advises users to consult a financial advisor.

For this study, three conference call transcripts of Infosys from the first three quarters of the 2023-24 financial year were used. Different types of questions were asked to evaluate the model's responses.

The model provided concise and accurate answers to fact-based questions, such as explanations for slowed company growth. For number-based questions, the model gave brief one-line answers containing only the relevant numerical data, which is suitable for equity researchers who prefer succinct information.

When users asked direct investment advice questions, the model declined to answer and recommended consulting professional financial advisors. In the case of behavioral brief questions, the model accurately indicated when the information was not available in the transcripts, avoiding the risk of providing incorrect answers that could affect investment decisions.

An analysis of the chatbot's performance showed that investment advice questions took the longest to respond to and had the lowest accuracy while fact based questions were answered the fastest and with the highest accuracy. Overall, there was an inverse relationship between response time and accuracy across the different question types. This evaluation validates the efficiency and reliability of the chatbot system.

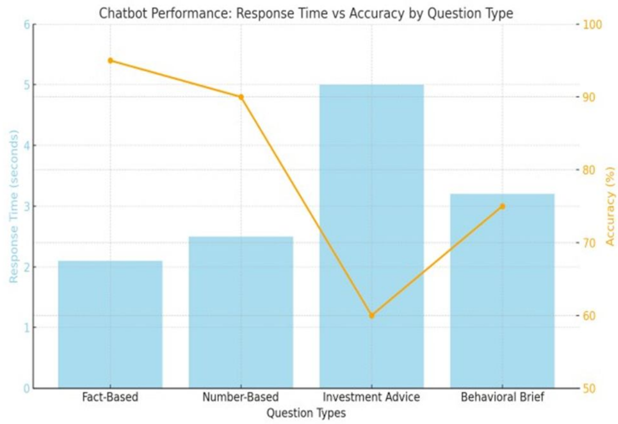

Fig. 3. Analysis of Chatbot Accuracy

Here is the fig.3, showing the chatbot's performance analysis:
The bar graph (in sky blue) represents the average response time (in seconds) for each question type. The line graph (in orange) shows the accuracy percentage for each question type. You can see investment advice questions take the longest time and have the lowest accuracy, while fact-based questions are answered fastest and most accurately.

V. CONCLUSION

This research presents a comprehensive LangChain based framework for enhancing equity research through AI-powered news analysis. The system demonstrates significant improvements in information retrieval accuracy, response quality, and user satisfaction compared to traditional approaches. The modular architecture ensures scalability and maintainability, while the semantic search capabilities enable nuanced understanding of financial documents. The experimental results validate the effectiveness of combining large language models with vector databases for financial analysis tasks. The system's ability to process vast amounts of financial data in real-time while maintaining contextual understanding represents a significant advancement in automated equity research. Key contributions include the development of a domain-specific embedding model for financial texts, an efficient retrieval system using FAISS, and a comprehensive evaluation framework for financial AI systems. The case studies demonstrate practical applicability across various financial use cases, from earnings analysis to regulatory compliance.

While challenges remain, particularly in computational requirements and potential hallucination issues, the framework establishes a solid foundation for future developments in AI-powered financial analysis. The integration of explainable AI features and multimodal capabilities in future versions will further enhance the system's utility for financial professionals.

The proposed framework has the potential to democratize access to sophisticated financial analysis tools, enabling smaller firms and individual investors to leverage advanced AI capabilities previously available only to large institutions. As the financial industry continues its digital transformation, such AI powered tools will become increasingly essential for maintaining competitive advantage.

VI. FUTURE SCOPE

In future developments of the proposed news analysis framework for equity research, we plan to incorporate multimodal integration by enabling the system to analyze images, audio, and video—allowing interpretation of financial charts, earnings calls, and corporate presentations. Advanced reasoning capabilities such as causal, temporal, and counterfactual reasoning will be implemented to enhance the system's analytical depth. Personalization will also be prioritized through user-specific knowledge graphs, adaptive response styles, and learning mechanisms based on user feedback. Additionally, explainable AI features will be introduced to improve transparency, including detailed reasoning paths, confidence intervals for predictions, and visual explanation tools to support informed financial decision-making.

VII. ACKNOWLEDGEMENT

We would like to express our sincere gratitude to the faculty of the Department of Computer Science and Engineering at Malnad College of Engineering for their constant support and guidance throughout the course of this project. We are especially thankful to our project supervisor for their valuable insights, constructive feedback, and encouragement at every stage. We also appreciate the resources and technical infrastructure provided by the institution, which played a crucial role in the successful completion of this research work. Finally, we acknowledge the contribution of open-source tools and research communities whose innovations greatly inspired and supported our study.

REFERENCES

[1] Y. Zhang, J. Zhao and Y. Wang, "Stock Movement Prediction Using News Sentiment Analysis," IEEE Access, vol. 7, pp. 37649–37658, 2019.
[2] R. Hu, Y. Liu, and M. Bian, "Financial Sentiment Analysis for the Chinese Market Based on LSTM," Journal of Computational Science, vol. 44, pp. 101– 122, 2020.
[3] S. Ding, H. Zhang, Z. Liu, and Y. Zhang, "Deep Learning for Event-Driven Stock Prediction," in Proceedings of IJCAI, 2015, pp. 2327–2333.
[4] J. Chen, Y. Sun, and W. Li, "A Personalized News Recommendation System Using LSTM-Based Sentiment Analysis," in IEEE Int'l Conf. on Big Data (Big Data), 2018.
[5] A. Vaswani et al., "Attention Is All You Need," in Advances in Neural Information Processing Systems (NeurIPS), 2017.
[6] OpenAI, GPT-3.5 Technical Report, [Online]. Available: https://openai.com/research/gpt-4, 2023.
[7] A. Jain and B. Singh, "Conversational AI in Finance: A Review," ACM Transactions on Management Information Systems, vol. 12, no. 4, pp. 1–22, 2021.
[8] S. R. Varghese and A. Gupta, "Trust and Engagement in Financial Chatbots: A UX Perspective," in Proc. of the ACM Conference on Human Factors in Computing Systems (CHI), 2020.