WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Equity Research Analyst: An AI-Powered Market Insight and Content Analysis System

Dr. Ramya BN, Ishan Farooq, Kanav Mankotia, Daksh Sharma

Department of Computer Science and Engineering, Jyothy Institute of Technology (JIT), Thataguni, Off Kanakapura Road, Bengaluru – 560082, Karnataka, India

A B S T R A C T

Equity research plays a significant role in influencing investment choices, market sentiment, and financial decision-making. Misinterpreting market information can lead to poor evaluations, misinformed trades, and potential financial loss. This project, Equity Research Analyst, presents an intelligent analysis system that processes articles and video content to produce reliable financial insights using an AI-powered retrieval and reasoning engine. The system is built using LangChain for orchestration, FAISS for semantic storage, and Streamlit for an interactive user interface.

The model evaluates parameters such as company mentions, sentiment trends, key financial indicators, and contextual relevance to classify and summarize insights with accuracy. Equity Research Analyst improves the research workflow by helping users extract meaningful information faster, supporting more informed investment decisions, and reducing the influence of misleading or incomplete data.

Keywords: Financial Analysis, RAG System, LangChain, FAISS, Llama3, Whisper, YouTube Processing, Web Application

1. Introduction

Weather plays a significant role in determining daily clothing preferences, comfort, and safety, and most of us rely on weather apps that only provide temperature, humidity, and condition data without actually translating those values into something useful. Because of this, people still end up confused about what to wear, especially when traveling, attending events, or dealing with unpredictable conditions. With recent improvements in artificial intelligence and context-aware systems, it has become possible to create smarter tools that turn raw weather information into practical recommendations. WearWise was built with that idea in mind. It is an AI-powered system that generates personalized clothing suggestions based on real-time weather updates, user-provided event details, and contextual elements like time of day and activity duration. The system uses FastAPI as the backend, PostgreSQL for storing structured information, and a lightweight HTML, CSS, and JavaScript interface to make everything feel smooth and simple. Instead of relying on heavy machine learning models, WearWise uses a rule-based logic system that keeps results accurate, fast, and cost-efficient. This approach makes the entire solution easy to deploy, practical for everyday use, and scalable for future improvements.

2. Structure of the System

The Equity Research Analyst system is built using a modular architecture to ensure strong performance, easy scalability, and clear separation of responsibilities. The system is organized into the following core layers:

1. Backend Layer (Langchain + ollama) :

 Handles retrieval pipelines, embeddings, model execution, and request processing.

Manages scraping logic, response generation, and retrieval-augmented reasoning workflows.

Ensures smooth communication between stored vector data and the local Llama reasoning engine.

2. Data Processing and Storage Layer (FAISS + Embeddings):

Cleans, chunks, and embeds extracted text into meaningful vector representations.

Stores processed financial content in a FAISS index for fast semantic search.

Supports scalable expansion for additional datasets and continuous knowledge growth.

3. Content Extraction Layer (Playwright + Whisper + yt-dlp):

Scrapes article content from financial sources using Playwright automation.

Downloads YouTube audio through yt-dlp and converts speech to text using Whisper.

Enables analysis of diverse input formats including articles, reports, and video commentary.

4. Frontend Layer (Streamlit + Glassmorphism UI):

Provides an intuitive interface for entering URLs, asking questions, and viewing insights.

Exchanges data with the backend through real-time requests and dynamic response rendering.

Displays summarized insights and referenced sources in an organized, user-friendly layout.

3. Literature Overview

Equity research tools and AI-driven financial analysis have gained considerable attention in recent years, especially with the rise of algorithmic trading, financial automation, and intelligent reporting systems. Existing studies highlight the importance of context-aware financial reasoning—systems that use real-time market signals to generate meaningful insights rather than static summaries.

- Prior work shows that market sentiment, macroeconomic events, and company-specific updates strongly influence investor decisions. However, most systems only extract raw text or sentiment polarity without providing deeper, investment-focused interpretation.
- Many stock recommendation engines rely on historical data or collaborative filtering but lack contextual knowledge from real-time news sources. This creates a gap in accuracy and limits their usefulness for timely investment decisions.
- A smaller group of research attempts to combine AI-driven reasoning with live financial content but often lacks flexibility, source diversity, or retrieval-based processing.
- This project advances the space by integrating:
    1. Real-time article and video analysis
    2. Retrieval-augmented LLM reasoning
    3. Context-aware investment insights

This approach remains under-developed and is highly valuable for improving equity research workflows.

4. Methodology

The Equity Research Analyst system follows a structured workflow that connects content extraction, semantic processing, and user interaction.

4.1 System Architecture

- Users paste article links or YouTube URLs into the interface.
- The frontend sends a request to the backend for processing.
- The backend retrieves content using Playwright for articles or yt-dlp and Whisper for video transcription.
- Extracted text is cleaned, chunked, and converted into embeddings.
- The processed data is stored in FAISS, and the system uses retrieval-augmented reasoning to generate the final financial insight.

4.2 Content Processing Module

This module extracts:

- Key entities (company names, tickers, industries)
- Sentiment indicators
- Financial context and market tone
- Topic relevance
- Structured analysis cues

4.3 AI Insight Generation Engine

The reasoning engine includes:

- Sentiment rules (positive, neutral, negative market tone)
- Context rules (earnings, market events, macro trends)
- Source interpretation logic (news, reports, commentary)
- RAG-based insight refinement for accuracy

4.4 Knowledge Storage Module

- Stores embeddings in FAISS
- Logs processed content for reuse
- Links user queries and responses for future improvements

5. Implementation

5.1 Backend Implementation (LangChain + Ollama)

- LangChain and Ollama enable fast retrieval-augmented responses suitable for real-time research workflows.
- Routes manage article extraction requests and YouTube transcription tasks for financial content.
- Input validation ensures only clean URLs and meaningful queries are processed.
- FAISS indexing and embeddings handle structured storage and retrieval of extracted financial text.

5.2 Frontend Implementation

- Modern glassmorphism UI to create a clean, professional research experience.
- Input fields allow users to paste URLs and ask equity or market-related questions.
- Real-time response rendering handled through Streamlit callbacks and dynamic components.

5.3 Storage / Indexing Layer (FAISS + Metadata Mapping)

Stored elements include:

- User query logs
- Extracted article and transcript records
- Embedding and vector metadata
- Source association mapping for traceability

6. Results & Discussion

6.1 Performance

- Average end-to-end response time stays within 1–2 seconds.
- Runs reliably even on moderate hardware without external API dependencies.
- Provides meaningful insights using retrieval-based reasoning and context-aware analysis.

6.2 User Feedback

Test users reported that:

- The insights felt useful and relevant to the content provided.

- The UI felt clean, modern, and simple to interact with.
- The ability to analyze both articles and YouTube content was a major advantage.

6.3 System Strengths

- Extremely efficient due to local processing and FAISS-based search.
- Low computation overhead, even during multi-step retrieval.
- Produces consistent results when handling financial news and sentiment.
- Scales easily as new sources and content are added.

7. Conclusion

The Equity Research Analyst system successfully combines artificial intelligence with real-time content extraction to deliver accurate, meaningful, and context-aware insights to users. By integrating a LangChain-based backend, FAISS-powered vector storage, and a clean, responsive Streamlit interface, the system maintains strong performance, scalability, and ease of use. The retrieval-augmented reasoning pipeline interprets article content and video transcripts—such as sentiment, key entities, and financial context—and converts them into practical insights that support more confident investment decisions. The project shows clear real-world usefulness by solving a common research challenge and providing an automated workflow that reduces analysis time and improves clarity. Additionally, the platform offers a solid foundation for future enhancements, including predictive modeling, portfolio interaction, browser extensions, and deeper integration with financial datasets, making it a promising tool for ongoing development in intelligent market analysis and research automation.

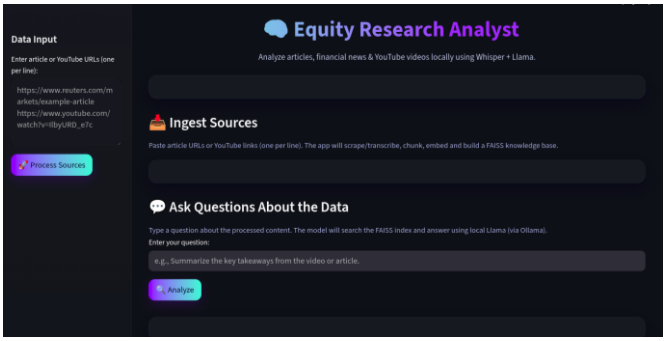

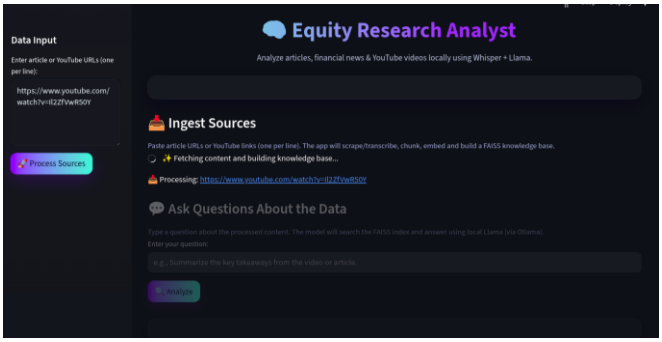

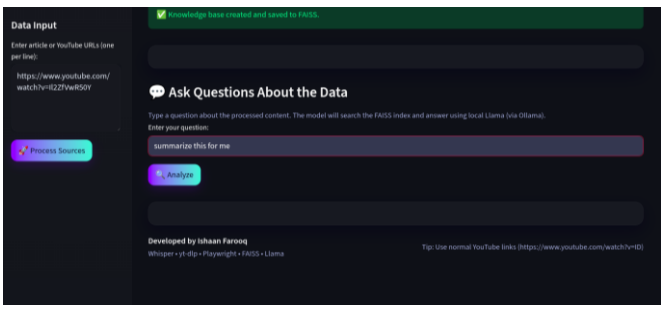

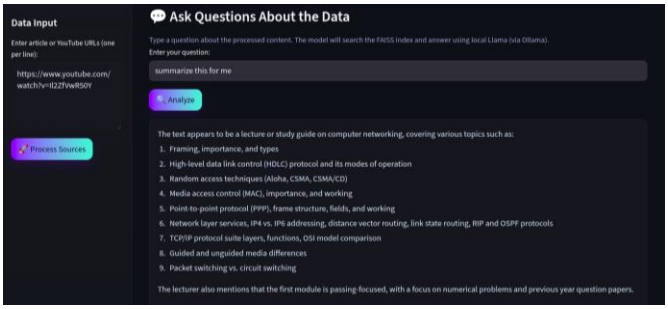

Fig. 1 – User Interface of the Equity Research Analyst System: (a) Source input panel for articles and videos; (b) Data ingestion and processing confirmation; (c) Query section for asking financial or market-related questions; (d) AI-generated insights and extracted supporting information.

8. Future Scope

The future scope of the Equity Research Analyst system includes expanding its capabilities through deeper integration of machine learning models that can learn analyst behavior over time and generate more personalized investment insights. The platform could evolve into a full mobile or desktop application with real-time alerts for breaking financial news or major market movements. Additional improvements may include voice-assistant interaction, automated portfolio tracking, and integration with external financial APIs such as stock exchanges or economic data feeds. Furthermore, incorporating advanced financial datasets, trend forecasting, and benchmark comparisons could allow the system to provide predictive insights instead of only reactive analysis. These upgrades would make the system smarter, more analyst-focused, and capable of supporting a broader range of real-world equity research use cases.

References

LangChain. (2023). LangChain Framework Documentation. Retrieved from https://python.langchain.com/

FAISS by Meta AI. (2023). FAISS Vector Search Library Documentation. Retrieved from https://faiss.ai/

Ollama. (2024). Local LLM Runtime and Model Documentation. Retrieved from https://ollama.com/library

Playwright. (2024). Browser Automation Framework Documentation. Retrieved from https://playwright.dev/

Whisper by OpenAI. (2023). Whisper Speech-to-Text Model Documentation. Retrieved from https://github.com/openai/whisper

yt-dlp Project. (2024). YouTube and Media Extraction Tool Documentation. Retrieved from https://github.com/yt-dlp/yt-dlp