



Conversational Stock Research Agent Using LLMs and MCP Data

M. R. C. V. Hruday Babu

M.S., Data Science

Department of Computer Science

Sri Venkateswara University, Tirupati, India

Email: mrcvhruday@gmail.com

M. Padmavathamma

Professor, Department of Computer Science

SVU College of CM & CS,

Sri Venkateswara University, Tirupati, India

Email: prof.padma@yahoo.com

Abstract

The financial research landscape has experienced an unprecedented rise in data complexity, driven by the rapid growth of real-time market streams, corporate disclosures, and global news sources. Traditional stock analysis requires extensive time, domain expertise, and manual effort, making it increasingly inaccessible for retail investors and operationally demanding for professional analysts. This research presents the design and implementation of a *Conversational Stock Research Agent* that leverages Large Language Models (LLMs) and MCP (Market and Corporate Performance) Data APIs to deliver automated, explainable, and interactive equity research assistance. The proposed system integrates structured MCP datasets with unstructured financial text using LangChain-based reasoning pipelines. It introduces three specialised Chain-of-Thought (CoT) modules: Data-CoT (quantitative analysis), Thesis-CoT (qualitative reasoning), and Risk-CoT (risk assessment). These modules enable contextual interpretation of numerical indicators and qualitative factors affecting stock performance. The backend, implemented using FastAPI and Flask, interfaces with the GPT API for inferential reasoning and orchestrates end-to-end data processing. Experimental results demonstrate that the agent can autonomously generate multi-layered financial reports—including growth metrics, risk exposure, and investment rationale—with accuracy comparable to professional analyst summaries. The system reduces report-generation time by more than 70 percent relative to manual research workflows. Future enhancements, including integration with real-time dashboards and portfolio analytics, position this framework as a scalable foundation for next-generation financial intelligence systems.

Keywords: Stock research, conversational AI, LangChain, MCP financial data, GPT API, risk analysis, FastAPI, MongoDB, Docker

I. Introduction

The stock market remains one of the most data-intensive environments, where timely and accurate insights can significantly influence investment outcomes. Investors traditionally rely on comprehensive research reports that synthesise historical performance metrics, financial ratios, qualitative commentary, and risk forecasts. However, manually analysing financial statements, balance sheets, and economic news is both labour-intensive and prone to human error. The exponential growth in financial data sources—including real-time market feeds, corporate filings, and macroeconomic news—has further amplified this complexity. Recent advancements in Large Language Models (LLMs), such as GPT-4, have introduced new possibilities for interpreting financial information through natural language reasoning. When combined with structured market datasets accessed through APIs, these models can automate substantial components of the equity research process, transforming static analytics into dynamic, conversational intelligence. The motivation for this project lies in bridging the gap between raw financial data and accessible, actionable insights. The Conversational Stock Research Agent enables users to pose queries such as, “What is the growth outlook of HDFC Bank?” or “Which companies in the IT sector exhibit the lowest risk-to-return ratio?” and receive human-like, data-grounded analyses. By unifying LLM-driven reasoning with MCP data retrieval, this system democratises financial research and empowers both professionals and retail investors with instant, context-rich intelligence.

II. Related Work / Literature Review

The automation of financial research has evolved through multiple stages, progressing from rule-based analytics to deep learning-driven insights. Early financial analysis platforms such as Yahoo Finance and Google Finance primarily offered static dashboards and descriptive historical visualisations, lacking interpretive or predictive capabilities. In the early 2010s, machine learning algorithms began to be applied to market prediction tasks, yet these models frequently suffered from limited interpretability and inconsistent generalisation. Between 2019 and 2021, transformer-based architectures improved contextual understanding of financial text, enabling models such as FinBERT to support sentiment analysis. However, such models remained task-specific and did not provide end-to-end, explainable reasoning. The development of LangChain and Retrieval-Augmented Generation (RAG) frameworks enabled the integration of external structured and unstructured knowledge sources into LLM-based workflows. These innovations made it possible to build domain-specific conversational agents capable of dynamically incorporating real-time financial data into natural language explanations.

Recent studies and systems demonstrate the progression of this field:

- **ChatGPT for Financial Insights (OpenAI, 2024)** showcased strong qualitative reasoning but lacked direct grounding in structured financial datasets.
- **FinGPT (Yang et al., 2024)** emphasised fine-tuning LLMs for market prediction but did not provide conversational explainability or transparent data sourcing.
- **BloombergGPT (Wu et al., 2023)** demonstrated the potential of large-scale financial language models but remained proprietary and resource-intensive.

Despite these advances, existing tools still face limitations in explainability, data transparency, and seamless integration of structured market data. The proposed research agent addresses these gaps by combining structured MCP data with LangChain’s modular reasoning pipelines, ensuring that every analytical conclusion is traceable, verifiable, and grounded in real data.

III. System Analysis / Problem Statement

A. Problem Identification

Traditional financial analysis tools are often either too complex for non-experts or too limited in their reasoning capabilities. Analysts must manually extract, interpret, and synthesise information from numerous sources, contributing to information overload, inconsistencies, and significant time expenditure. Retail investors, in particular, lack access to the sophisticated, end-to-end analytical insights available to professional research teams.

Key Challenges Identified:

1. **Manual report generation is time-consuming**, often requiring several hours per company analysis.
2. **Unstructured financial text**—such as news articles, disclosures, and commentary—is underutilised due to the difficulty of contextual interpretation.
3. **Existing tools lack explainable, query-driven financial reasoning**, limiting transparency and user trust.
4. **Integration of heterogeneous data sources** (e.g., MCP datasets, market APIs, unstructured text) remains technically challenging and error-prone.

B. Objectives

The primary objectives of the proposed system are:

- To build a conversational AI system capable of real-time financial reasoning.
- To integrate structured (MCP data) and unstructured (news, disclosures, reports) financial information using LangChain workflows.
- To automate the generation of multi-section equity research reports (e.g., company overview, financial analysis, valuation, and risk assessment).
- To ensure transparency and explainability in all AI-generated insights, with clear traceability to underlying data sources.

C. System Requirements

1) Functional Requirements

The system shall:

- Support query handling and data retrieval through a FastAPI-based backend.
- Ingest financial data from MCP (Market and Corporate Performance) APIs, including balance sheets, KPIs, and other fundamental indicators.
- Perform LLM-based reasoning using the GPT API integrated via LangChain.
- Use MongoDB for persistent storage of chat sessions, embeddings, and generated reports.
- Provide a web-based user interface (UI) that supports natural language interaction for equity research queries.

2) Non-Functional Requirements

The system shall:

- Achieve an average response time of ≤ 5 seconds per query under normal load.

- Ensure secure API authentication and encryption for all external and internal data exchanges.
- Support horizontal scalability and containerised deployment using Docker.

IV. System Design / Methodology

A. Architectural Overview

The overall architecture follows a three-tier design:

1. Presentation Layer (Flask UI)

A Flask-based web interface, implemented using HTML, CSS, and JavaScript, provides an interactive chat-like environment for users. It allows investors and analysts to enter natural language queries and view structured, explanatory responses, including charts and tabular outputs.

2. Application Layer (FastAPI Backend)

The backend, built with FastAPI, serves as the core application layer. It:

- Orchestrates LangChain workflows and chain executions.
- Manages communication with external MCP data APIs.
- Interfaces with the GPT API for LLM-based reasoning.
- Implements prompt engineering and context construction using LangChain prompt templates.
- Routes requests asynchronously between the UI, data layer, and reasoning modules.

3. Data Layer

The data layer consists of:

- **MCP APIs** provide real-time and historical market and corporate performance data (e.g., stock prices, financial statements, key performance indicators).
- **MongoDB** for session storage, embeddings, user queries, and generated analytical outputs.

The system architecture integrates multiple functional layers, each employing specialised technologies to ensure scalability, modularity, and efficient communication. The Data Retrieval Module leverages MCP Financial Data APIs and Python-based HTTP clients to fetch real-time structured data. The Prompt Engineering Layer uses LangChain to build context-aware prompts for the LLM dynamically. The LLM Reasoning Module, powered by GPT APIs, performs multi-step inferential analysis, combining quantitative signals with qualitative insights. A Response Formatting layer structures the generated outputs for readability and downstream visualisation. This design transforms traditional one-dimensional dashboards into an interactive research assistant. By combining structured MCP data with unstructured textual analysis and LLM reasoning, the system reduces cognitive load, enhances interpretability, and delivers timely, conversational financial intelligence.

B. LangChain-Based Reasoning

The system employs three specialised LangChain-based reasoning modules (chains), each designed for a specific aspect of equity research:

- **Data-CoT (Quantitative Chain-of-Thought)**

Focuses on numerical and quantitative indicators such as earnings per share (EPS), revenue growth, profit margins, valuation multiples, and key financial ratios. It interprets trends and compares metrics across time or against sector benchmarks.

- **Thesis-CoT (Qualitative Chain-of-Thought)**

Evaluates qualitative aspects, including the business model, competitive positioning, industry trends, management quality, and sentiment derived from news, reports, and disclosures. It helps formulate an investment thesis grounded in both data and narrative context.

- **Risk-CoT (Risk Chain-of-Thought)**

Assesses risk factors such as price volatility, leverage and debt exposure, governance concerns, regulatory risks, and sector-specific vulnerabilities. It highlights downside scenarios and risk-to-return considerations.

The outputs of these three chains are merged into a coherent, multi-level response. The final answer is formatted for human readability, typically structured into sections such as “Company Overview,” “Financial Health,” “Growth Outlook,” “Risk Analysis,” and “Investment View.”

C. Data Pipeline

The end-to-end data pipeline for the Conversational Stock Research Agent consists of the following stages:

1. **Data Retrieval from MCP APIs**

The system fetches financial data via MCP APIs, including balance sheets, income statements, cash flow data, KPIs, and market metrics.

2. **Tabular-to-JSON Conversion**

Raw tabular outputs from MCP APIs are cleaned, normalised, and converted into structured JSON formats. This ensures consistent representation of fields across different companies and time periods.

3. **Preprocessing and Normalisation**

- Handle missing values, inconsistent units, and currency normalisation where required.
- Standardise key indicators (e.g., EPS, revenue growth, P/E ratio) into a uniform schema.

4. **Integration with LangChain Pipelines**

The structured JSON data is fed into LangChain modules:

- **Data-CoT** for quantitative reasoning.
- **Thesis-CoT** for qualitative and contextual analysis.
- **Risk-CoT** for risk profiling and scenario analysis.

5. **LLM-Based Analytical Narrative Generation**

Using the GPT API, the system generates a structured analytical narrative that explains the numbers, highlights trends, and formulates an investment rationale. This narrative is grounded in the MCP data and contextual prompts constructed by LangChain.

6. **Response Formatting and Visualisation**

The analytical outputs are formatted using templating tools such as Python Jinja2. Visual elements (charts, graphs) are produced using libraries such as Chart.js or Matplotlib to illustrate trends in revenue, profitability, valuation, and risk metrics.

7. **Session Persistence in MongoDB**


All user queries, intermediate reasoning states (where appropriate), and final generated reports are stored in MongoDB. This persistence:

- Enables session continuity and conversational context.
- Supports auditability and traceability of AI-generated insights.
- Allows users to revisit prior analyses and compare historical recommendations.

Through this pipeline, the system delivers end-to-end, explainable financial intelligence—starting from raw MCP data and user queries and culminating in interactive, conversational research reports.

D. Methodological Flow

The methodological workflow integrates data retrieval, preprocessing, LLM-based reasoning, and multi-format output generation. The end-to-end flow is as follows:



Architectural Overview

1. User submits a natural language query via the Flask UI.
2. FastAPI backend receives the query and extracts key entities (company, period, financial metric).
3. MCP API is invoked to retrieve structured financial data.
4. Data is cleaned, normalised, and converted into JSON representations.
5. LangChain pipelines (Data-CoT, Thesis-CoT, Risk-CoT) process the structured and unstructured inputs.

6. GPT API generates a consolidated analytical narrative.
7. MongoDB stores session history, embeddings, and generated outputs.
8. Flask UI renders the results with charts and summary insights.
9. Users may export the generated content as a PDF report.

V. Implementation

A. Tools and Technologies

Component	Technology
Frontend	Flask, HTML, CSS, Chart.js
Backend	FastAPI, Python
AI Framework	LangChain, GPT API
Database	MongoDB
Data Source	MCP Financial APIs
Deployment	Docker, Uvicorn

B. Backend Logic

The FastAPI server orchestrates the end-to-end reasoning workflow. Each user request triggers the following sequence:

1. **MCP API Fetch:** Retrieve financial statements, KPIs, valuation metrics, or risk indicators.
2. **Prompt Construction:** Combine MCP data, company metadata, and user intent into a context-rich prompt using LangChain templates.
3. **LLM Invocation:** Pass the structured prompt to the GPT API for inferential reasoning.
4. **Storage:** Persist user queries, responses, and intermediate reasoning states in MongoDB for traceability and session continuity.

C. Flask Frontend

The Flask-based UI provides:

- A natural language input box for research queries.
- Display of AI-generated summaries, financial interpretations, and interactive charts.
- An export feature for generating PDF versions of research reports.

D. Example Query Flow

User Query:

“Explain HDFC Bank’s risk exposure for FY2024.”

System-Generated Response (Excerpt):

“HDFC Bank exhibits stable earnings with moderate credit exposure. The NPA ratio improved by 0.4% year-over-year, and capital adequacy remains above regulatory thresholds. Primary risks are associated with retail lending concentration.”

Additional analysis:

- The loan portfolio remains diversified across retail, corporate, and SME segments, supporting consistent earnings performance.
- Interest rate sensitivity remains moderate due to prudent asset–liability management.
- Macro risk exposure is limited by the bank’s strong capital buffers and high-quality collateral.

VI. Results and Discussion

A. Experimental Setup

- Analysis conducted for **50 companies** across BSE and NSE.
- MCP API data supplemented with synthetic financial statements for controlled testing.
- User interaction and system performance were evaluated across **25 real-world query scenarios**.

B. Performance Metrics

Metric	Result
Average Response Time	3.2 seconds
Quantitative Accuracy	94.7%
Risk Assessment Accuracy	91.3%
User Satisfaction (Survey)	4.6 / 5

C. Observations

- Modularisation through LangChain significantly improved interpretability and reduced output ambiguity.
- MCP integration enhanced the factual grounding and credibility of generated insights.
- Qualitative reasoning capabilities made the system accessible and valuable for non-technical investors.
- Transformer-based architectures, when combined with structured datasets, proved effective in automating complex equity research workflows.

D. Comparative Analysis

Approach	Manual Analysis	Proposed Agent
Avg. Report Generation Time	30 minutes	3 minutes
Data Integration	Manual	Automated
Explainability	Limited	High
User Accessibility	Analysts only	All users

VII. Conclusion and Future Work

The Conversational Stock Research Agent demonstrates how the integration of LLMs with structured market data can automate and enhance the financial research process. The system streamlines report



generation, reduces dependence on manual data interpretation, and delivers scalable explainability through modular Chain-of-Thought reasoning.

Future Enhancements:

1. Incorporation of predictive modelling for earnings and revenue forecasts.
2. Expansion to cover mutual funds, ETFs, derivatives, and macroeconomic indicators.
3. Deployment of mobile and cloud-native versions for broader accessibility.
4. Integration of reinforcement learning mechanisms to adapt to user feedback and improve recommendation quality.

This research contributes meaningfully to the emerging field of AI-driven financial intelligence and demonstrates a blueprint for scalable, transparent, and interactive equity research systems.

References

- [1]. Wu, S. et al., "BloombergGPT: A Large Language Model for Finance," Bloomberg Research, 2023.
- [2]. Yang, W. et al., "FinGPT: Open Financial Large Language Models," arXiv:2306.06031, 2024.
- [3]. OpenAI, "GPT-4 Technical Report," arXiv:2303.08774, 2023.
- [4]. Wei, J. et al., "Chain-of-Thought Reasoning in Large Language Models," ICLR, 2023.
- [5]. Capgemini, "Generative AI in Financial Services," Industry Report, 2024.
- [6]. McKinsey Global Institute, "The State of AI in 2024," McKinsey Report, 2024.
- [7]. Chiarello, F. et al., "Financial Analytics Using LLM-Oriented Architectures," FGCS Journal, 2025.
- [8]. LangChain Documentation, "Prompt Templates and RAG Framework," LangChain.org, 2025.
- [9]. Dataiku, "LLM-Powered Finance Automation," Dataiku Blog, 2025.
- [10]. Patel, V., Das, G., "AI in Market Intelligence," IEEE Transactions on Engineering Management, 2023.