LTIMindtree

WHITEPAPER

Empowering Fund Managers with Multi-Agent AI Financial Technology

LTIMindtree

Abstract

Today, fund managers face numerous challenges and spend considerable effort in synthesizing and interpreting information from various sources such as research data, market data, financial statements, and fund performance metrics. They need to distinguish between noise in the data and meaningful signals to make informed investment decisions. However, with the increasing volume and variety of data, it can be difficult to identify the most relevant and actionable insights.

With new AI financial technology and Large Language Models (LLMs), fund managers can analyze large amounts of structured and unstructured data, identify patterns, and provide insights that may not be apparent to humans. However, the LLMs with traditional Retrieval-Augmented Generation (RAG) models, which are effective for general document retrieval and synthesis, fall short when addressing the complex needs of fund managers. Fund managers require not just the retrieval of relevant information but also the real-time interpretation of complex, multi-modal financial data, including structured elements like tables, charts, and performance metrics, along with unstructured data such as market commentary.

Traditional RAG models excel at retrieving static documents but lack the depth and flexibility to combine and analyze these diverse data sources effectively. Additionally, these RAG systems cannot generate visual aids such as charts, tables, or plots, which are crucial for data-driven financial decision-making.

This whitepaper introduces a multi-agentic AI financial technology for fund managers to address challenges with traditional RAG. It is aimed at handling complex user queries related to financial decision-making. This financial automation system features multiple agents, including the Fact Finder, Fund Navigator, Reviewer, Researcher, and Chart Wizard, all of which collaborate to retrieve, analyze, and present multi-modal financial data.

The system's Fact Finder is particularly adept at extracting relevant information from a variety of sources, such as tables, charts, market commentaries, and internal knowledge bases like market insights, fund reports, and government policies before passing it to the Large Language Model (LLM) for response generation. Meanwhile, other agents, such as the Fund Navigator and Reviewer, validate the information and ensure its accuracy and relevance. This multi-agent approach provides a robust framework for delivering comprehensive financial insights.

Table of Contents

Introduction

The financial sector operates in an era of rapid data generation, where vast volumes of market updates, policy shifts, performance reports, and visual data need integration for decision-making. Traditional tools and methods often struggle to keep pace with this data influx, leading to missed opportunities and suboptimal decisions. To address these challenges, a multi-agentic AI financial system was developed that transforms data retrieval and analysis for fund managers. The system combines LLMs with a specialized multi-agent framework to provide:

- **Multi-agentic Collaboration:** Specialized agents work together to handle user queries effectively.
- **Fact Finder Expertise:** Extracts and synthesizes information from tables, charts, and market commentaries to provide actionable insights.
- **Enhanced Depth with LLMs:** Large Language Models' integration enriches the clarity and depth of responses.
- **Comprehensive Query Coverage:** Agents like Fund Navigator, Reviewer, Researcher, and Chart Wizard address all aspects of each query.

This innovation enables fund managers to make faster, more accurate decisions by automating complex data analysis. Financial system automation helps streamline the entire decision-making process, providing users with more reliable results.

Multi-Agents Overview

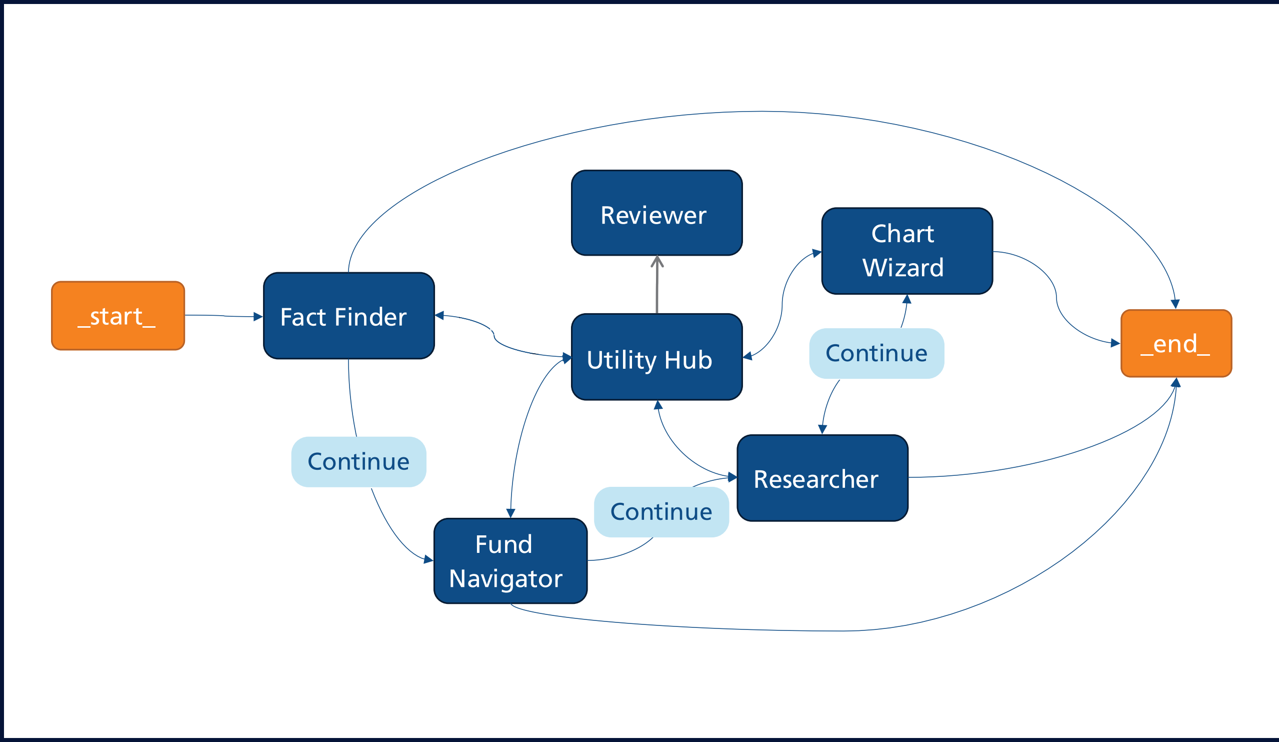

Figure 1. Flow of an Agent call

To fully appreciate the capabilities of this AI financial technology, it's essential to explore the roles of each agent within the system. These agents work collaboratively to process diverse user queries related to financial decision-making:

- **Fact Finder (Agent 1):** This agent serves as the first point of interaction. It retrieves relevant contexts from a vector database built from internal knowledge sources like market commentary, government policies, market insights, fund reports, and market update tables or charts.
  The Fact Finder is designed to parse structured data (e.g., tables, charts) and unstructured data (e.g., textual commentary) to provide holistic financial context.

- **Fund Navigator (Agent 2):** The Fund Navigator searches the fund repository for relevant funds based on the user's query and context provided by the Fact Finder. It ensures that recommended funds align with the user's investment preferences and current market trends.

Multi-Agents Overview

- **Reviewer (Agent 3):** The Reviewer works in parallel with the Fact Finder and Fund Navigator to ensure that all retrieved data is accurate, relevant, and up to date before it is presented to the user.
- **Researcher (Agent 4):** The Researcher conducts deeper analysis and provides more nuanced insights based on the Fact Finder's and Fund Navigator's outputs. The Researcher is responsible for ensuring that emerging trends or potential opportunities are included in the final output. When required, this system can generate charts, tables, and visual representations, adding a multi-modal dimension. Moreover, it is intelligent in attaching personalized recommendations for investment ideas, guaranteeing that fund managers receive actionable insights tailored to specific financial strategies.
- **Chart Wizard (Agent 5):** This agent converts complex financial data into visual representations such as charts, graphs, and tables. The Chart Wizard works closely with the Researcher to produce meaningful visual data to accompany the system's textual outputs.

The advanced multi-agent AI financial system offers nuanced, real-time analysis. It handles diverse data types and presents them through clear visualizations. In contrast, traditional RAG systems rely on singular data types and static updates, limiting their scope. The multi-agent system ensures accuracy, relevance, and timeliness, making it ideal for complex decision-making.

Now that we have a clear understanding of each agent's role, let's explore how they collaborate to handle complex queries.

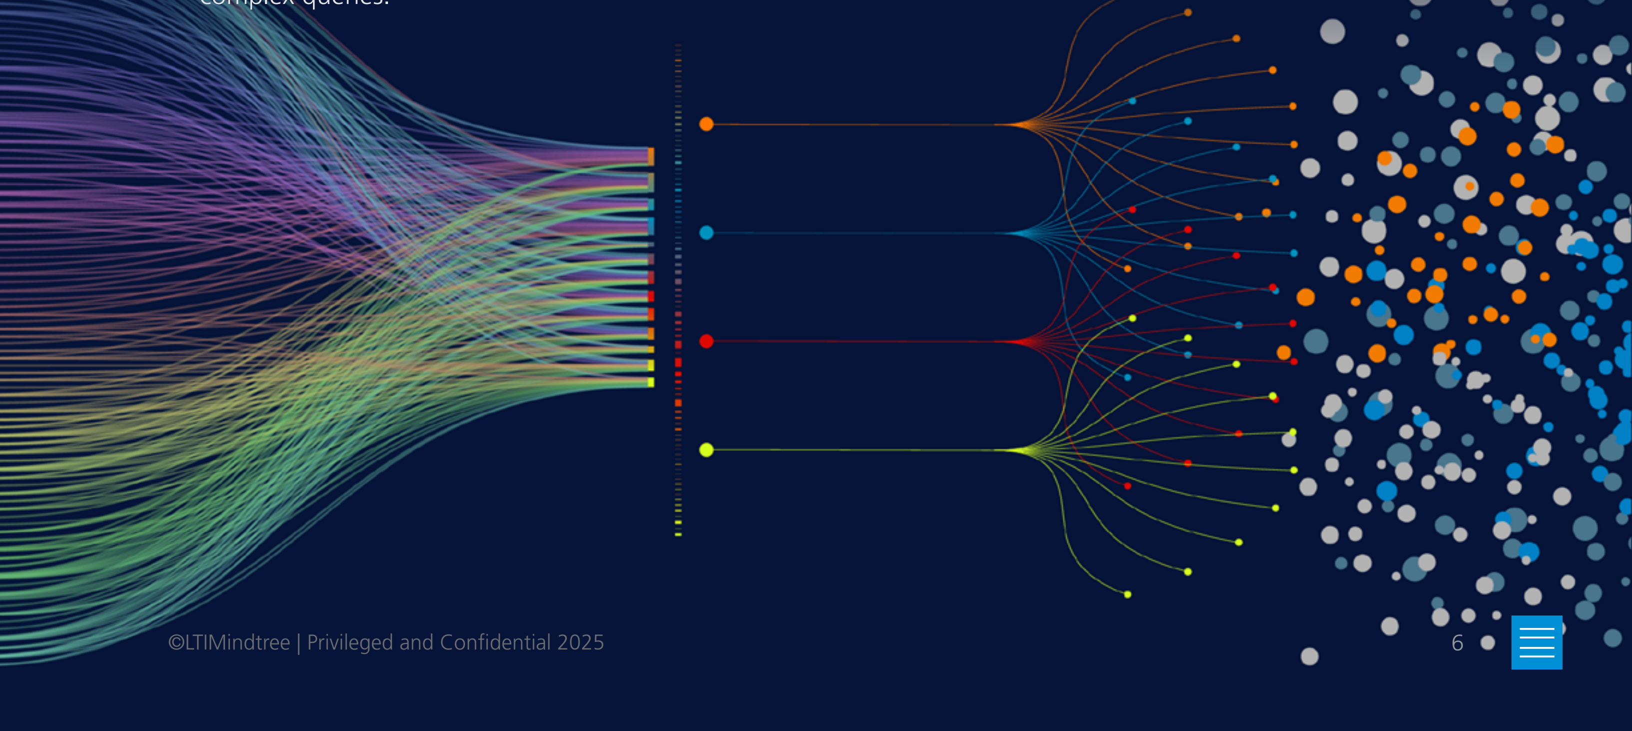

Workflow of the Multi Agentic System

The system's workflow ensures seamless interaction among agents, enabling accurate, timely insights for fund managers:

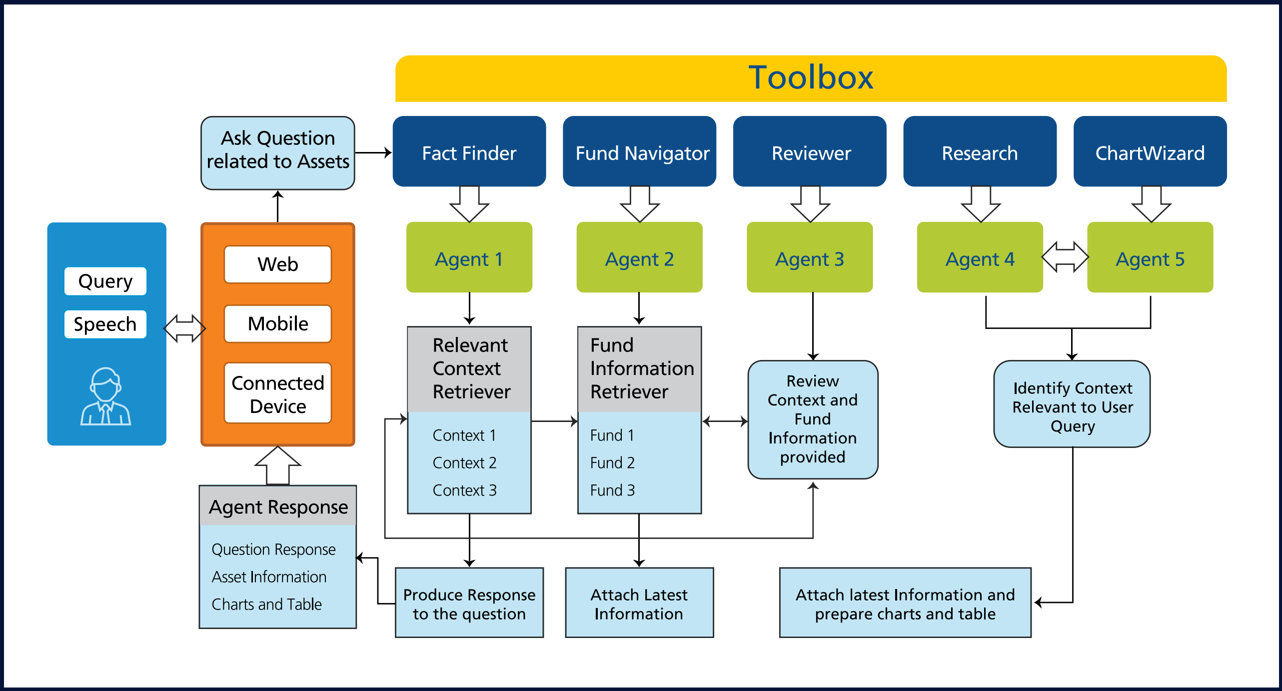

Figure 2. Architecture for Agentic Workflow

1 Query Submission and Context Retrieval

The user submits a query through the web or mobile interface. This query could involve topics like market performance, fund comparisons, or investment strategies. The Fact Finder then begins by extracting relevant context from the vector database, which contains internal knowledge sources such as market commentary, government policies, market update tables, and market insights.

Workflow of the Multi Agentic System

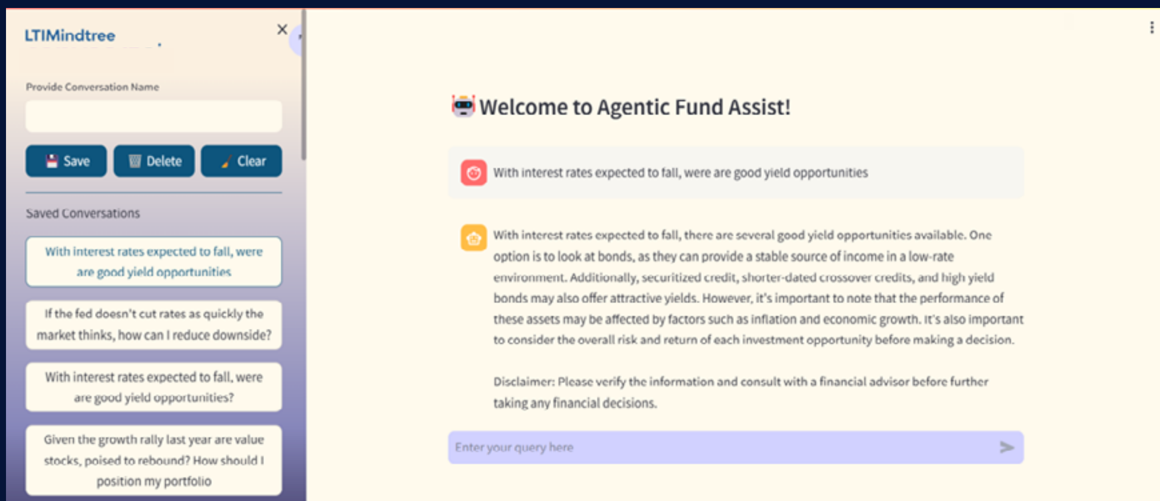

Figure 3. A WebUI for Chat Interface of Fund Assistant

In addition to extracting data from unstructured texts, the Fact Finder is designed to parse and interpret structured data from tables and charts. For instance, if the user query pertains to sector performance, the Fact Finder might extract figures from a table detailing sector-wise fund performance. Similarly, if a market commentary includes graphical data on market trends, the Fact Finder integrates those insights into its context retrieval process.

Then, the Fact Finder forwards the gathered context to the LLM for response generation. The LLM synthesizes the retrieved information into a coherent response tailored to the user’s query.

2 Parallel Processing by Fund Navigator and Reviewer

While the LLM processes the information, the Fund Navigator searches the fund repository to identify relevant funds. These funds align with the context provided by the Fact Finder and the LLM-generated response. For example, if the query asks for sustainable energy funds, the Fund Navigator will return a list of top-performing funds in this sector, based on the insights retrieved.

Simultaneously, the Reviewer monitors outputs from both the Fact Finder and Fund Navigator. The Reviewer ensures the data used is accurate, relevant, and consistent with the latest market updates. It acts as a quality control mechanism, ensuring only precise and reliable information reaches the user.

Workflow of the Multi Agentic System

This parallel processing speeds up the system and ensures decision-making is based on current, accurate data. It enhances efficiency by:

- Increasing productivity, enabling more work at the same time.
- Improving user experience by providing quicker responses and reducing wait times.
- Enhancing reliability by cross-verifying data through multiple processes, leading to more accurate outcomes.
- Facilitating scalability, enabling systems to handle more complex tasks or a larger volume of tasks without a linear increase in processing time.

These benefits make the system more robust and efficient, capable of handling complex queries and delivering high-quality responses, essential for fields such as financial services, where timely and accurate information is crucial.

3 Researcher and Data Visualization

The Researcher investigates complex queries by utilizing additional data sources and conducting deeper analysis. This enriches the output generated by the LLM and Fund Navigator. For instance, the Researcher might dive deeper into market trends or provide historical performance data to support the recommendations made by the system.

The Chart Wizard collaborates with the Researcher to generate visual aids such as charts, graphs, and tables. These visualizations help users interpret the data. For example, if the query involves sector performance, the Chart Wizard might generate a comparative chart displaying the performance of different sectors over a specified time.

4 Final Output

Once all agents have completed their tasks, the system compiles the data and presents the final output to the user. This includes the LLM-generated response, fund recommendations from the Fund Navigator, and visual data representations (if necessary) from the Chart Wizard.

Technical Foundation

The effectiveness of this financial system automation relies on its robust technical underpinnings, which include:

Vector Database for Context Retrieval

The Fact Finder retrieves data from a vector database that stores structured and unstructured financial information. This database is curated from sources such as market commentary, market insights, fund reports, and government policy documents. It also includes tables and charts, which the Fact Finder parses to extract relevant numerical insights.

Large Language Model (LLM) Integration

Once the context is retrieved, the LLM generates a response based on the query. The integration of structured (from tables and charts) and unstructured data (market commentary, reports) allows LLM to provide more nuanced and comprehensive responses to user queries.

Fund Repository and Fund Selection

The Fund Navigator searches for a comprehensive fund repository to find funds that match the user's preferences and the context retrieved by the Fact Finder. This repository includes data on historical performance, risk factors, and sector allocations.

Multi-agent Parallel Processing

The Reviewer and Researcher agents work in parallel, ensuring the retrieved data is accurate, relevant, and supported by deeper insights. This parallel processing enables the system to manage complex queries in real-time, providing users with comprehensive, reliable answers.

Workflow of the Multi-Agentic System in Amazon Bedrock

The system's workflow ensures seamless interaction among agents, enabling accurate, timely insights for fund managers:

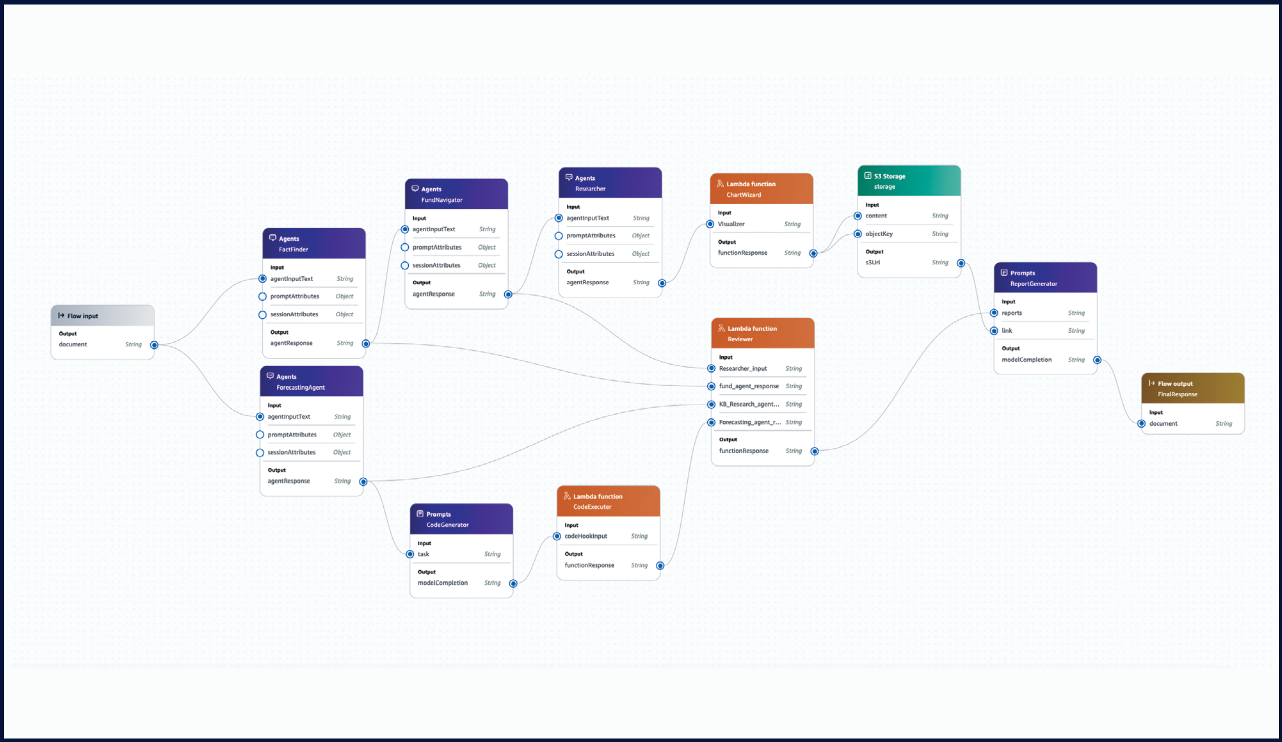

Figure 4. Making Agentic Workflow in AWS Flows

Future Directions

AI financial technology for financial decision-making could evolve significantly through the integration of traditional ML models, AI-driven automation, and data from marketplaces.

Integration with Traditional ML Models:

The system would gain the ability to forecast market trends and investment outcomes, helping fund managers make proactive decisions based on anticipated shifts in sectors, fund performance, and the impact of major events.

AI-driven Automation:

This automates routine tasks, such as portfolio rebalancing, alerting fund managers of key changes, and generating compliance reports, thereby enhancing efficiency and reducing human error.

Data Market Place:

Data from marketplaces would provide access to proprietary information, expanding the system's capabilities and integrating directly into client-facing platforms for enriched financial services.

Altogether, these advancements would transform the system into a powerful, adaptive tool that enhances fund managers' ability to make timely, data-driven decisions in a dynamic financial landscape.

Use Case Example: Query on Renewable Energy Funds

A user submits a query asking for the best-performing renewable energy funds. The Fact Finder extracts relevant data from market commentary and performance reports, pulling key metrics from tables and charts that detail the performance of various renewable energy funds. The LLM generates a response by summarizing the findings.

Simultaneously, the Fund Navigator identifies top-performing renewable energy funds from the fund repository, providing specific fund recommendations based on the context retrieved by the Fact Finder. The Reviewer ensures the accuracy of both the context and fund recommendations.

The Researcher enriches the output with a detailed analysis of future trends in the renewable energy sector, while the Chart Wizard generates a visual comparison of the top funds' performance over the past year.

The final output includes the LLM-generated response, fund recommendations, and a performance chart, providing the user with a well-rounded answer to their query.

Conclusion

The multi-agentic AI financial system improves financial decision-making by integrating diverse data sources such as tables, charts, and market commentary. It addresses the limitations of traditional RAG methods. The system's multi-agent architecture allows parallel processing, ensuring users receive accurate, comprehensive, and actionable insights. With the inclusion of agents like Fact Finder, even complex queries are handled effectively. This reduces the manual effort fund managers typically spend analyzing data and improving decision-making.

To know more about how we are solving complex AI problems, please visit our website at LTIMindtree.

About Authors

__Neeraj Verma__ is an Associate Principal Data Scientist with over 17 years of experience in designing and deploying scalable AI and Machine Learning solutions. With expertise spanning Business Intelligence, Deep Learning, and cloud platforms like AWS, he helps organizations build efficient, cloud-native architectures that unlock data-driven growth.

__Sateesh Gottumukkala__ is a Senior Director in Enterprise AI, with more than 15 years of experience in advanced analytics, machine learning, and artificial intelligence. He has a proven track record of designing end-to-end AI solutions and implementing MLOps to solve complex business challenges across industries.

__Ashish Varerkar__ is a transformational AI leader with over 28 years of experience in AI, Cloud, Digital Transformation, and Consulting. As the head of LTIMindtree's Enterprise AI practice, he guides organizations in adopting AI and Generative AI technologies with a focus on achieving tangible business outcomes. A certified Enterprise Architect, Ashish combines strategic insight with hands-on expertise to help clients navigate their AI journeys effectively.

__Bhadresh__ is a seasoned Senior Data Scientist specializing in Natural Language Processing (NLP), Machine Learning, and Data Science. He focuses on developing and deploying AI solutions that address real-world challenges, leveraging open-source tools to foster innovation and drive measurable impact.

Citations

1. Retrieval-Augmented Generation : Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks

2. FinCon LLM Multi-Agent System: FinCon: A Synthesized LLM Multi-Agent System with Conceptual Verbal Reinforcement for Enhanced Financial Decision Making

3. Multi-Agent Reinforcement Learning: Multi-agent reinforcement learning and financial optimization systems in portfolio management

4. AWS Flows for Agentic Flow: AWS FLows for Building Agentic Flow

5. AWS Bedrock Flow Tracing: AWS Bedrock Flow Tracing

LTIMindtree

Let's get to the future, faster. Together.