SETSCI Conference Proceedings 24, 17-22, 2025

International Conference on Open Finance

August 28, 2025 | Western New England University, Springfield, MA, USA

ICOF 2025

https://doi.org/10.36287/setsci.24.2.017

AICrypto-Assistant: A Multi-Agent LLM Platform for Democratizing Crypto-Asset Analysis

Mimoza Dimodugno[1] , Mehdi Mammadov[1]

[1]*Northeastern University, Boston, MA, USA (m.dimodugno@northeastern.edu)*

Abstract – Market democratization through conversational AI represents a critical frontier in financial technology, yet cryptocurrency analysis remains dominated by institutional players with sophisticated toolchains. We introduce AICrypto-Assistant, the first open-source platform to integrate multi-agent large language model orchestration with executable technical analysis in a real-time conversational interface. Our system employs supervisor that intelligently routes natural language queries to specialized agents: a Market Analyst executing Python-based indicators (moving averages, volume-price relationships, money flow metrics), a News Researcher synthesizing market developments through filtered sentiment analysis, and a Knowledge Management component delivering contextual information via retrieval-augmented generation. Unlike existing black-box solutions, every analytical step remains transparent and auditable, addressing critical barriers in accessibility, information overload, and methodological opacity that exclude retail participants. Controlled evaluation with 20 participants over four weeks demonstrated remarkable effectiveness: 11% collective portfolio growth, 95% user satisfaction, and 60% achieving positive returns despite market volatility. The modular architecture enables seamless integration of additional indicators without core system modifications, supporting long-term extensibility while advancing Open Finance principles of transparent, inclusive financial intelligence.

Keywords – Cryptocurrency, FinTech, ExplainableAI, Multi-Agent Systems, Open Finance, Retail Investing

I. Introduction

The cryptocurrency market has evolved from Bitcoin's debut into a three-trillion-dollar financial ecosystem [1], yet this growth has exacerbated analytical disparities between institutional and retail participants. Professional trading desks synthesize technical indicators, blockchain forensics, macroeconomic data, and sentiment analytics. Retail investors typically rely on simplified dashboards, social media sentiment, or intuition—approaches that increase vulnerability to market volatility and behavioral biases.

This analytical divide contradicts the democratization principles underlying both Decentralized Finance and Open Finance, which advocates transparent, equitable access to financial tools and information. While public blockchains provide open transaction data, the interpretive capabilities required to extract actionable insights remain concentrated among specialists with advanced technical knowledge or expensive analytical software.

The convergence of large language models (LLMs) and high-throughput financial APIs presents an opportunity to bridge this gap. Modern LLMs demonstrate exceptional capability in translating technical concepts into accessible explanations, while APIs like Binance now provide millisecond-latency market data. However, existing cryptocurrency analysis tools typically suffer from two limitations: proprietary systems that obscure analytical methodologies, or superficial chatbots that provide definitions without executing quantitative analysis.

AICrypto-Assistant addresses these shortcomings through a multi-agent architecture that combines natural language interaction with transparent, executable analysis modules. By separating language processing from computational logic and maintaining full open-source accessibility, the system delivers professional-grade analytics while preserving the methodological transparency essential to Open Finance principles.

II. Problem Statement & Motivation

We identify four structural barriers preventing retail traders from accessing professional-grade cryptocurrency analysis capabilities, creating persistent information asymmetries that contradict Open Finance principles.

A. Analytic Accessibility Deficit

Professional cryptocurrency traders employ sophisticated analytical frameworks integrating multiple technical indicators, on-chain metrics, and sentiment signals. Retail investors encounter fragmented resources—isolated charting platforms, social media channels, and subscription services—each providing limited perspective. Even when platforms automate indicator calculations, they assume users understand the interpretation of crossover events, divergence patterns, and trend confirmations. This knowledge prerequisite creates cognitive barriers that discourage participation or lead to emotion-driven trading decisions rather than systematic analysis.

B. Information Retrieval Overload

Cryptocurrency markets exhibit heightened sensitivity to external events compared to traditional assets. Regulatory announcements, protocol updates, security incidents, and macroeconomic developments can trigger significant price movements within minutes. Relevant information must be identified across thousands of sources—news platforms, developer forums, social media, and regulatory publications. The volume and velocity of this information flow exceed human processing capacity, making comprehensive monitoring impractical without computational assistance. Existing automated solutions either lack comprehension or generate excessive noise without adequate filtering.

C. Adaptability Constraints

Cryptocurrency participants range from professional traders requiring granular parameter control to newcomers seeking guided explanations of fundamental concepts. Most platforms force a binary choice: advanced tools with complex interfaces that overwhelm beginners, or simplified alternatives that limit analytical depth. This segregation creates distinct markets for expert versus novice tools, with limited pathways for progression between sophistication levels.

D. Open Finance Imperative

The cryptocurrency ecosystem emerged from principles of transparency, verifiability, and open participation—values essential to Open Finance. Yet many contemporary AI-powered trading assistants operate as proprietary systems with undisclosed methodologies, creating a trust paradox where users rely on opaque systems to interpret supposedly transparent markets [3]. This contradiction undermines decentralized finance's premise that effective market participation should not require trusting centralized authorities.

These barriers collectively maintain information asymmetries that limit retail investor effectiveness while contradicting the democratization goals underlying both cryptocurrency adoption and Open Finance initiatives. AICrypto-Assistant addresses these challenges through transparent, modular architecture that preserves analytical rigor while eliminating technical barriers to access.

III. Methodology & System Architecture

AICrypto-Assistant employs a modular, multi-agent architecture that delivers cryptocurrency analysis, news research, and educational content through an interactive web interface. The system combines a Flask-based backend with LangGraph-orchestrated agents to provide transparent, extensible analytical capabilities.

A. System Architecture Overview

The platform implements a "router + specialists" pattern using LangGraph, a framework for orchestrating interactions between large language models and external computational tools. This architecture separates query understanding, domain-specific analysis, and result presentation while maintaining user experience coherence. Figure 1 illustrates the complete system architecture.

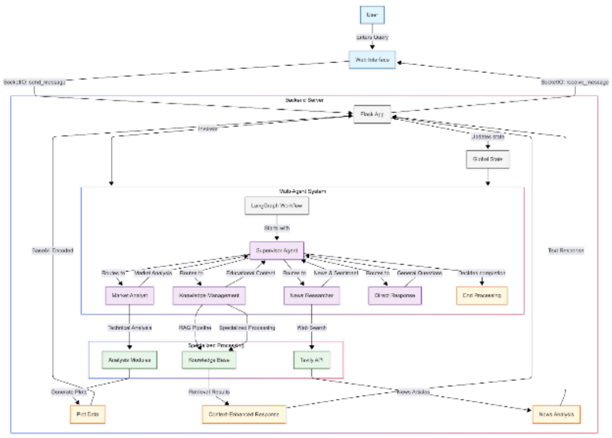

Figure 1 Query Flow Architecture

The modular design features four specialized agents: the Knowledge Management agent manages informational queries through retrieval-augmented generation, the Market Analyst executes technical indicator analysis, the News Researcher processes market sentiment, and the Direct Response agent manages general conversation. This separation enables independent development and maintenance of system components while preserving unified interaction patterns.

B. Query Processing Workflow

The temporal sequence of processing a user query is illustrated in Figure 2, highlighting the step-by-step flow of information through the system. When a user submits a query, the web interface transmits it to the Flask backend via SocketIO. The backend appends the query to the conversation history and invokes the LangGraph workflow.

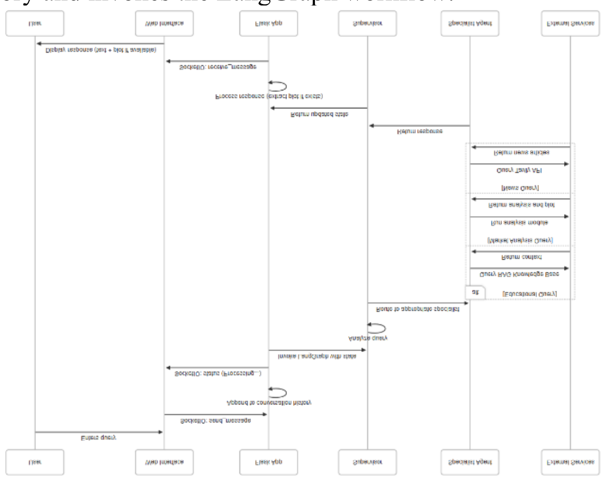

Figure 2 Query Processing Sequence

The workflow begins with the Supervisor agent, which analyzes the query and routes it to the appropriate specialist. Depending on the query type, different processing paths are followed:

- Educational queries are routed to the Knowledge Management agent, which leverages a Retrieval Augmented Generation (RAG) pipeline to access a knowledge base of cryptocurrency concepts.
- Market analysis queries are directed to the Market Analyst, which selects and executes appropriate analysis modules to generate interpretations and visualizations.

- News-related queries are managed by the News Researcher, which interfaces with the Tavily API to fetch and analyze recent cryptocurrency news.

After processing, specialist agents return responses to the Supervisor, which may request additional information or conclude the interaction. The final response, potentially including visualizations, is returned to the Flask backend, which formats it for display and transmits it to the web interface via SocketIO.

C. Data Flow Analysis

Figure 3 provides a detailed view of how data flows through the system components, highlighting the specific data transformations at each step of the process.

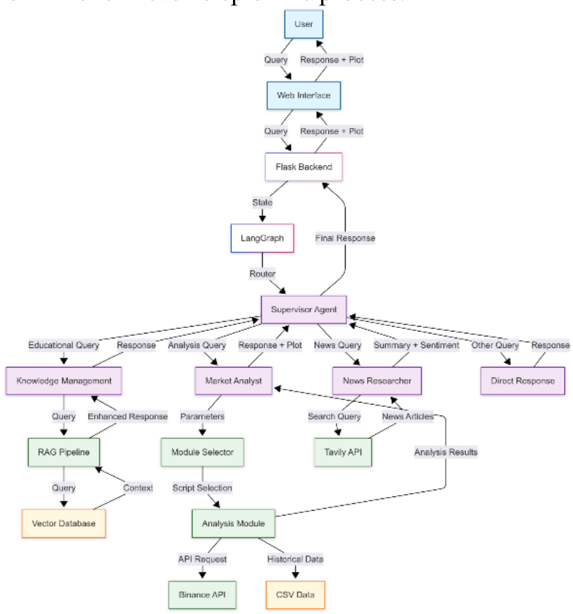

Figure 3 Data Flow Diagram

In the data flow, user queries are transformed through multiple processing stages:

1. **Query Routing**: The Supervisor agent performs initial query analysis to determine the appropriate processing path.
2. **Specialized Processing**:
   - Educational queries trigger a RAG pipeline that retrieves relevant information from a vector database before generating a contextually enhanced response.
   - Analysis queries undergo parameter extraction and module selection before being processed by specialized analysis modules. These modules may fetch data from external APIs (such as Binance Api) or local CSV data, generate interpretations, and produce visualizations.
   - News queries are processed through the Tavily API to retrieve relevant articles, which are then summarized and analyzed for sentiment.
3. **Response Integration**: All specialist responses flow back through the Supervisor to the Flask backend, which manages the extraction and formatting of textual content and visualizations.

The data flow architecture ensures efficient processing of different query types while maintaining a consistent interface for the user. This design allows for future extension with additional specialist agents or data sources without requiring changes to the core system architecture.

IV. Related Work

Recent advancements in LLMs have spurred significant interest in their application to financial analysis, particularly in the volatile domain of cryptocurrency trading. This section reviews key works on multi-agent LLM systems, AI-driven financial tools, and explainable AI in finance, positioning our AICrypto-Assistant platform within the current research landscape.

A. LLM-Based Multi-Agent Systems in Finance

The use of multi-agent systems powered by LLMs has emerged as a promising approach for complex financial tasks. For instance, a framework for automated cryptocurrency portfolio management leverages multiple LLM agents to integrate multi-modal data, such as market trends and news sentiment, to enhance decision-making [10]. This system emphasizes explainability, aligning with our platform's commitment to transparency in line with Open Finance principles. Similarly, TradingAgents proposes a multi-agent LLM framework inspired by trading firms, featuring specialized agents for fundamental, sentiment, and technical analysis [11]. While these works share our multi-agent approach, AICrypto-Assistant uniquely focuses on democratizing access for retail investors through an open-source platform with real-time conversational capabilities.

B. AI-Driven Financial Analysis Tools

Several studies have explored LLM-based agents for financial analysis beyond cryptocurrencies. FinSphere introduces a conversational stock analysis agent equipped with quantitative tools and real-time data [2]. MarketSenseAI 2.0 enhances stock analysis through LLM agents, focusing on specialized tasks like technical indicator computation [7]. These systems provide valuable benchmarks for AICrypto-Assistant, though our platform's emphasis on cryptocurrency-specific challenges, such as information overload and adaptability, distinguishes it. Additionally, a study on numerical understanding in LLM-based trading agents highlights limitations in handling quantitative data, an area where our platform's modular design and specialized Market Analyst agent aim to improve performance [6].

C. Explainable AI and Open Finance

Transparency is a critical concern in AI-driven financial systems, particularly in the context of Open Finance. A comprehensive study on multi-agent LLMs in traditional and decentralized finance demonstrates their potential to address inefficiencies and regulatory requirements, achieving a 3.8% increase in portfolio returns [9]. This aligns with our platform's goal of providing auditable and transparent methodologies. Furthermore, surveys on LLM-based multi-agent systems offer insights into collaboration mechanisms and challenges, 1 such as role flipping and memory management, which inform our design choices [4][8].

D. Positioning AICrypto-Assistant

While existing works advance the application of LLMs in finance, AICryptoAssistant stands out by integrating educational components (via the Knowledge Management agent), real-time interaction through Flask-based backend, and a commitment to Open Finance principles. Unlike proprietary systems, our open-source approach ensures accessibility, and our preliminary evaluation results, showing an 11% portfolio growth, suggest competitive performance. Future work can build on these foundations to address remaining challenges, such as enhancing numerical accuracy and scalability.

V. Preliminary Evaluation / Results

To assess the effectiveness of AICrypto-Assistant in addressing the identified barriers to cryptocurrency analysis, we conducted a preliminary user study with 20 volunteer participants. This section presents the methodology and results of this evaluation, focusing on both quantitative performance metrics and qualitative user feedback.

A. Evaluation Methodology

We recruited 20 participants with varying levels of cryptocurrency trading experience, ranging from complete beginners to experienced traders. Each participant was provided with:

- A test account on Binance with 100 USDT (totaling 2000 USDT across all participants)
- Access to the AICrypto-Assistant system
- Freedom to make trading decisions based on the system's analysis and recommendations
- A post-study questionnaire to capture qualitative feedback

The evaluation period lasted for four weeks, during which participants could use the system at their discretion while making trading decisions on their test accounts. We tracked both financial performance and user interaction patterns throughout the study period.

B. Quantitative Performance Metrics

a. Financial Performance

The primary quantitative metric for evaluation was the change in portfolio value over the study period. As shown in Figure 4, the collective portfolio value across all participants increased from the initial 2000 USDT to 2220 USDT by the end of the evaluation period, representing an 11% overall return.

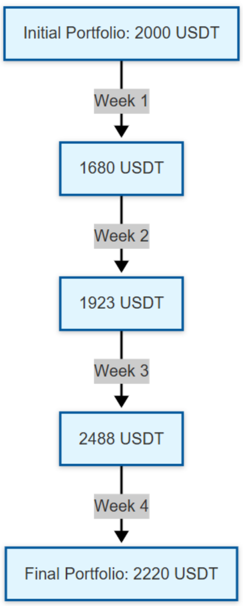

Figure 4 Portfolio Value Growth During Evaluation Period

Participants demonstrated varying levels of performance, as shown in Figure 5. A majority (60%, n=12) achieved positive returns, while a small subset experienced a minor loss (n=6) or broke even (n=2). The highest individual return was 18.5%, while the largest loss was 5.2%.

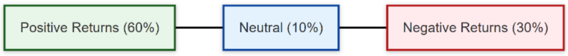

Figure 5 Distribution of Individual Returns

b. Strategy Effectiveness

We also analyzed the effectiveness of different analytical strategies provided by AICrypto-Assistant. Figure 6 shows average returns achieved when participants followed recommendations from each specialist agent.

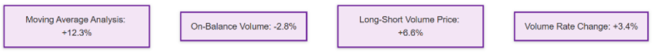

Figure 6 Average Returns by Analysis Strategy

The data indicates that Moving Average Analysis and Long-Short Volume Price analysis yielded the highest average returns among the technical indicators employed by AICrypto-Assistant. This suggests that these strategies were particularly effective during the evaluation period's market conditions.

c. System Usage Patterns

Throughout the evaluation period, we tracked how participants interacted with different components of the system. Figure 7 illustrates the distribution of queries across the specialist agents.

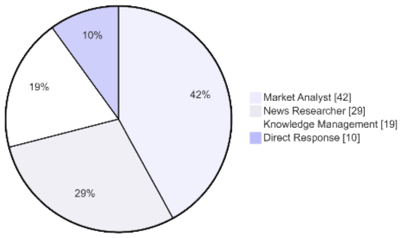

Figure 7 Distribution of User Queries by Agent Type

The Market Analyst agent received the highest proportion of queries (42.3%), followed by the News Researcher (28.7%), indicating that participants prioritized technical analysis and news sentiment when making trading decisions.

C. Qualitative User Feedback

At the conclusion of the study, participants completed a detailed questionnaire about their experience with AICrypto-Assistant. The results were overwhelmingly positive, with 19 out of 20 participants (95%) reporting satisfaction with the system. Key feedback themes are summarized below:

Figure 8 User Feedback Summary

The most frequently cited benefits of the system included:

- **Accessibility of complex concepts** - Participants with limited prior knowledge appreciated the system's ability to explain technical indicators in straightforward language.
- **Integrated information** - Users valued the combination of technical analysis, news research, and educational content within a single interface.
- **Adaptive interactions** - Both beginners and experienced traders found the system helpful, with

beginners citing the educational components and experienced traders valuing the analytical depth.

- **Transparency in analysis** - Users reported greater trust in recommendations due to the system's explicit explanations of analysis methodology.

One participant noted:
"*What impressed me most was how the system didn't just tell me what to do - it explained why a particular indicator was signaling a trend and helped me understand the underlying principles. This made me more confident in my decisions.*"

D. Impact on Open Finance Barriers

The evaluation results demonstrate AICrypto-Assistant's effectiveness in addressing the four critical barriers identified in our problem statement:

- **Analytic Accessibility Deficit** - The 95% positive feedback rate, particularly the high score for educational value (4.8/5), suggests the system successfully bridges the knowledge gap between professional and retail investors. The natural language interface allows users to access sophisticated analytical tools without requiring prior technical expertise.
- **Information Retrieval Overload** - The substantial usage of the News Researcher agent (28.7% of queries) indicates participants found value in the system's ability to filter and contextualize market news, reducing information overload while maintaining awareness of relevant developments.
- **Adaptability Constraints** - The positive feedback across participants of varying experience levels confirms the system's success in adapting to different user needs through the same interface. The natural language interaction model eliminates the traditional trade-off between simplicity and analytical power.
- **Open Finance Imperative** - The system's transparent explanation of analytical methodologies, coupled with high trust ratings, suggests progress toward maintaining open finance principles while leveraging advanced AI capabilities.

E. Limitations and Future Work

While the preliminary evaluation shows promising results, we acknowledge several limitations that will be addressed in future work:

Limited sample size - A larger and more diverse participant pool would provide more robust insights into the system's effectiveness across different user segments.

Short evaluation period - Extending the evaluation period would allow assessment of the system's performance across different market conditions.

Baseline comparison - Future studies will include direct comparisons with existing cryptocurrency analysis platforms to quantify relative advantages.

Feature expansion - Based on user feedback, we plan to expand the analysis modules to include additional technical indicators and on-chain metrics.

The next phase of development will focus on scaling the system architecture to accommodate a larger user base while maintaining **response** quality and speed. We also plan to implement a continuous learning mechanism that improves analytical accuracy based on historical performance.

VI. Discussion

AICrypto-Assistant stands out as a transformative tool for retail investors by breaking down structural barriers in cryptocurrency trading. A key strength of AICrypto-Assistant lies in its ability to address the four structural barriers outlined in our problem statement. The 95% user satisfaction rate, coupled with qualitative feedback praising the accessibility of complex concepts, demonstrates the platform’s success in mitigating the analytic accessibility deficit. By employing natural language processing and retrieval-augmented generation, the system empowers users—regardless of prior expertise—to engage with sophisticated tools traditionally reserved for institutional traders. The heavy reliance on the Market Analyst (42.3% of queries) and News Researcher (28.7% of queries) agents further illustrates how participants leveraged these features to inform their trading strategies, effectively reducing information retrieval overload by delivering concise, contextualized insights.

The platform’s adaptability to diverse users’ needs also merits attention. Positive feedback from both novice and experienced traders highlights its capacity to overcome adaptability constraints, offering an interface that balances simplicity with analytical depth. This versatility eliminates the traditional divide between beginner-friendly and expert-oriented tools, fostering a more inclusive analytical environment. Moreover, the Open Finance imperative is upheld through the platform’s transparent, auditable methodologies and open-source architecture. Users’ reported trust in the system’s recommendations—attributed to clear explanations of analytical processes, which reinforces the importance of transparency as a cornerstone of financial democratization.

The system's emphasis on educational components—evidenced by the Knowledge Management agent's substantial usage and high educational value ratings (4.8/5)—addresses a critical market infrastructure gap. Rather than simply automating trading decisions, the platform builds analytical literacy among participants. This educational approach may contribute to more informed market participation and reduced susceptibility to manipulation or panic-driven trading behaviours.

While the evaluation highlights these strengths, it also reveals areas for growth. The small sample size and short testing period suggests that broader, longer-term studies could better assess the platform’s impact on user education and retention. Future iterations might incorporate personalized guidance, such as tailored learning paths, to further enhance accessibility and support diverse user needs. These improvements could amplify the platform’s ability to empower retail investors over time.

As the cryptocurrency market continues to evolve, tools like AICrypto-Assistant could play a pivotal role in shaping a more equitable and efficient financial landscape.

Conclusion

AICrypto-Assistant demonstrates how multi-agent AI systems can democratize access to sophisticated cryptocurrency analysis tools. By combining natural language processing with transparent computational modules, the platform addresses the analytical accessibility deficit faced by retail traders without sacrificing methodological rigor. Our preliminary evaluation shows promising results in both user satisfaction and trading performance, suggesting that well-designed AI assistants can effectively bridge knowledge gaps in complex financial domains.

The open-source, modular architecture established in this work provides a foundation for future extensions while maintaining alignment with Open Finance principles of transparency and accessibility. As cryptocurrency markets continue to evolve, tools that reduce information asymmetry between professional and retail participants will play an increasingly important role in fostering market efficiency and inclusivity. Future work will focus on expanding analytical capabilities, enhancing personalization, and testing performance across diverse market conditions. Through continued development, AICrypto-Assistant aims to empower retail investors with professional-grade analytical tools in an accessible, educational format.

REFERENCES

[1] "Live cryptocurrency charts & market data | CoinMarketCap," *CoinMarketCap*. https://coinmarketcap.com/charts/

[2] S. Han, Q. Zhang, Y. Yao, W. Jin, and Z. Xu, "LLM Multi-Agent Systems: Challenges and open Problems," *arXiv.org*, Feb. 05, 2024. https://arxiv.org/abs/2402.03578

[3] I. Okpala, A. Golgoon, and A. R. Kannan, "Agentic AI Systems Applied to tasks in Financial Services: Modeling and model risk management crews," *arXiv.org*, Feb. 08, 2025. https://arxiv.org/abs/2502.05439

[4] S. Han, J. Zhang, Y. Shen, K. Yan, and H. Li, "FinSphere, a Real-Time stock analysis agent powered by Instruction-Tuned LLMs and domain tools," *arXiv.org*, Jan. 08, 2025. https://arxiv.org/abs/2501.12399

[5] H. Ding, Y. Li, J. Wang, and H. Chen, "Large language model agent in Financial Trading: A survey," *arXiv.org*, Jul. 26, 2024. https://arxiv.org/abs/2408.06361

[6] T. Ma, J. Du, W. Huang, W. Wang, L. Xie, X. Zhong, and J.T. Zhou, "LLM Knows Geometry Better than Algebra: Numerical Understanding of LLM-Based Agents in A Trading Arena," *arXiv.org*, Feb. 25, 2025. https://arxiv.org/abs/2502.17967

[7] G. Fatouros, K. Metaxas, J. Soldatos, and M. Karathanassis, "MarketSenseAI 2.0: Enhancing Stock Analysis through LLM Agents," *arXiv.org*, Feb. 01, 2025. https://arxiv.org/abs/2502.00415

[8] K.-T. Tran, D. Dao, M.-D. Nguyen, Q.-V. Pham, B. O'Sullivan, and H. D. Nguyen, "Multi-Agent Collaboration Mechanisms: A survey of LLMs," *arXiv.org*, Jan. 10, 2025. https://arxiv.org/abs/2501.06322

[9] A. De La Cruz, "Multi-Agent Large Language Models for Traditional Finance and Decentralized Finance", *Journal of Industrial Engineering & Applied Science*, vol. 3, no. 1, pp. 10–19, Feb. 2025. https://www.suaspress.org/ojs/index.php/JIEAS/article/view/v3n1a02

[10] Y. Xiao, E. Sun, D. Luo, and W. Wang, "TradingAgents: Multi-Agents LLM Financial Trading Framework," *arXiv.org*, Dec. 28, 2024. https://arxiv.org/abs/2412.20138

[11] Y. Xiao, E. Sun, D. Luo, and W. Wang, "TradingAgents: Multi-Agents LLM Financial Trading Framework," *arXiv.org*, Dec. 28, 2024. https://arxiv.org/abs/2412.20138