Journal of Industrial Engineering and Applied Science
Journal Home: http://jieas.suaspress.org/ | CODEN: JIEAAE
Vol. 3, No. 1, 2025 | ISSN 3005-6071 (Print) | ISSN 3005-608X (Online)

Multi-Agent Large Language Models for Traditional Finance and Decentralized Finance

DE LA CRUZ, Antio [1*]

[1] Bloomberg Bazil AI Lab, Brazil

***** DE LA CRUZ, Antio is the corresponding author, E-mail: antoniodelacruz777@outlook.com***

Abstract: This paper investigates the application of multi-agent large language models (LLMs) in traditional finance (TradFi) and decentralized finance (DeFi), with a focus on addressing challenges such as inefficiencies, security vulnerabilities, and regulatory requirements. We propose a conceptual framework for implementing multi-agent LLMs, which combines the advanced reasoning and language capabilities of LLMs with the collaborative nature of multi-agent systems. Through detailed case studies, we demonstrate the effectiveness of this framework in portfolio management, fraud detection, smart contract optimization, and regulatory compliance. Empirical data from these case studies show significant improvements in performance metrics such as accuracy, speed, and cost efficiency compared to traditional approaches. For instance, in portfolio management, multi-agent LLMs achieved a 3.8% increase in average returns and a 0.6 improvement in the Sharpe ratio. In DeFi, smart contract optimization reduced vulnerabilities by 66%, while decentralized lending protocols saw a 17% increase in liquidity utilization. This paper concludes with recommendations on how to improve effectiveness of LLMs and outlines future research directions, including the integration of quantum computing and the development of explainable AI techniques.

Keywords: Multi-agent, LLM, TradFi, DeFi, Agent, Decision Making Process.

Disciplines: Artificial Intelligence Technology. **Subjects:** Machine Learning.

DOI: https://doi.org/10.70393/6a69656173.323634 **ARK:** https://n2t.net/ark:/40704/JIEAS.v3n1a02

1 INTRODUCTION

The financial landscape has been driven by the convergence of artificial intelligence (AI) and blockchain. Traditional finance, characterized by centralized institutions such as banks, stock exchanges, and regulatory bodies, has long been the backbone of global economic systems. However, the emergence of decentralized finance: a blockchain-based ecosystem that eliminates intermediaries through smart contracts and peer-to-peer networks, has introduced a paradigm shift in how financial services are designed, delivered, and consumed.

The rapid advancement of large language models, such as OpenAI's GPT series, Google's Bard, and Meta's LLaMA. These models have demonstrated remarkable capabilities in natural language processing, reasoning, and decision-making, making them invaluable tools for automating complex tasks across industries. In finance, LLMs are increasingly being used for applications such as sentiment analysis, fraud detection, and customer support. However, the complexity and interconnectedness of modern financial systems, particularly in DeFi, demand more sophisticated AI solutions that can operate collaboratively and adaptively.

This is where multi-agent systems come into play. Multi-agent systems (MAS) are computational frameworks in which multiple intelligent agents interact to achieve shared or individual goals. When combined with LLMs, multi-agent systems can enable decentralized, collaborative decision-making, making them particularly well-suited for addressing the challenges of both TradFi and DeFi. For instance, in TradFi, multi-agent LLMs could optimize portfolio management by simulating diverse market scenarios, while in DeFi, they could enhance the security and efficiency of smart contracts by autonomously identifying vulnerabilities and proposing fixes.

This paper seeks to explore the potential of multi-agent LLMs in addressing the challenges of TradFi and DeFi. Specifically, the research aims to:

1. Investigate how multi-agent LLMs can enhance decision-making, automation, and interoperability in financial systems.
2. Identify key use cases for multi-agent LLMs in TradFi and DeFi, such as portfolio optimization, fraud detection, smart contract auditing, and regulatory compliance.
3. Analyze the technical, ethical, and regulatory challenges associated with deploying multi-agent LLMs in finance.
4. Propose a conceptual framework for designing and

Journal of Industrial Engineering and Applied Science
Journal Home: http://jieas.suaspress.org/ | CODEN: JIEAAE
Vol. 3, No. 1, 2025 | ISSN 3005-6071 (Print) | ISSN 3005-608X (Online)

implementing multi-agent LLMs in financial ecosystems.

2 LITERATURE REVIEW

2.1 TRADTIONAL FINANCE

Traditional finance refers to the conventional financial system comprising centralized institutions such as banks, stock exchanges, insurance companies, and regulatory bodies. These institutions have long served as intermediaries, facilitating transactions, managing risk, and ensuring compliance with legal and regulatory frameworks. Key challenges include inefficiencies in transaction processing, high operational costs, and limited accessibility for underserved populations.

The integration of technology into TradFi has been a gradual process, with early innovations such as electronic trading platforms and algorithmic trading systems paving the way for more advanced applications of artificial intelligence (AI). For instance, AI-powered tools are now widely used for credit scoring, fraud detection, and portfolio management [1]. LSTM based machine learning has shown superior results in the work of [2]. Our research is inspired by the LSTM framework and benchmark sets created by [2]. Despite these advancements, the adoption of cutting-edge technologies in TradFi remains constrained by legacy systems, regulatory hurdles, and resistance to change.

2.2 DECENTRALIZED FINANCE

Decentralized finance represents a radical departure from TradFi, leveraging blockchain technology to create an open, permissionless, and transparent financial ecosystem. DeFi eliminates intermediaries by using smart contracts—self-executing programs that run on blockchain networks such as Ethereum—to automate financial services like lending, borrowing, and trading [3]. The DeFi ecosystem has grown exponentially in recent years, with total value locked (TVL) in DeFi protocols surpassing $100 billion at its peak.

Despite its potential, DeFi presents several challenges in implementation. These include smart contract vulnerabilities, scalability issues, and regulatory uncertainty. Moreover, the lack of centralized oversight in DeFi creates unique risks, such as market manipulation and liquidity fragmentation [4]. These challenges highlight the need for innovative solutions that can enhance the security, efficiency, and interoperability of DeFi systems.

2.3 LARGE LANGUAGE MODELS

Large language models are a class of AI models trained on vast amounts of text data to understand and generate human-like language. Models such as OpenAI's GPT series, Google's Bard, Deepseek and Meta's LLaMA have demonstrated remarkable capabilities in natural language processing (NLP), reasoning, and decision-making [5]. In finance, LLMs are increasingly being used for applications such as sentiment analysis, customer support, and financial forecasting.

However, the application of LLMs in finance is not without challenges. These include issues related to data privacy, model interpretability, and the potential for bias in decision-making [6]. Moreover, single-agent LLMs are often limited in their ability to handle complex, multi-stakeholder environments, such as those found in financial ecosystems.

Our optimization-based approach started from ideas from transport industry research like [7]. We have studied previous work on [8], [9], [10] and [11]. This is a sample text, please replace it manually. This is a sample text, please replace it manually. This is a sample text, please replace it manually. This is a sample text, please replace it manually.

We have leveraged on the research of [12], [13] , [14] and [15]. Labeling and semantic encoder has been studied in [16] and [17]. The intelligent approach for texture detection has been studied in [18], [19] and [20].

Previous effort on detecting risk in digital asset transactions have shown promising results in [21] and [22]. In addition to the deep reinforcement learning approach from [23], [24] and [25], we leverage on the compositional continual learning for language models from [26] and [27]. The optimization part of such training can be borrowed from the weighted fuzzy rough sets-based tri-training example [28]. This can be further optimized by attribute reduction as illustrated in [29]. On the application side there are many examples of applying parameters or input attribute optimization like [30], [31] and [32].

Market sentiment is always an important element to our multi-agent LLM model. Past effort on market sentiment has been summarized in [33] and [34]. Input sentiments need to be processed and applied adaptively, which are explained in works of [35], [36] and [37]. Past empirical analysis also showed promising results from these literature references.

[38] and [39] demonstrate the power of LLM in financial forecasting. We have also found the hybrid method with LLM are particularly useful with reduced pre-processing times as shown in [40]. [41] has provided an alternative method in gold price prediction. [42], [43] and [44] enlighten us with pose estimation and distributed networking. Last but not least, we leverage the corner attention network from [45] and sparse fusion policy from [46] for our numerical analysis.

While significant progress has been made in the fields of TradFi, DeFi, LLMs, and MAS, there are notable gaps in literature. First, there is limited research on the integration of LLMs with multi-agent systems, particularly in the context of financial applications. Second, the potential of multi-agent LLMs to address the challenges of TradFi and DeFi remains underexplored. For instance, how can multi-agent LLMs enhance the security and efficiency of smart contracts in DeFi? How can they improve decision-making and regulatory compliance in TradFi?

Finally, there is a lack of comprehensive frameworks for designing and implementing multi-agent LLMs in financial ecosystems. Addressing these gaps is critical to unlocking the full potential of multi-agent LLMs in transforming traditional and decentralized finance.

3 MULTI-AGENT LLMS: CONCEPTUAL FRAMEWORK

3.1 WHAT IS MULTI-AGENT LLMS?

Multi-agent large language models are advanced systems that combine the reasoning and language capabilities of LLMs with the collaborative and adaptive nature of multi-agent systems. In a multi-agent LLM framework, multiple intelligent agents powered by an LLM work together to achieve shared or individual goals. This framework is like an well-trained colony of bees. These agents can communicate, negotiate, and coordinate their actions, making them particularly well-suited for complex, dynamic environments like financial ecosystems.

Unlike single-agent LLMs, which operate in isolation, multi-agent LLMs leverage the collective intelligence of multiple agents to solve problems that are beyond the scope of any single agent. For example, in a financial context, one agent might analyze market trends, another might assess risk, and a third might execute trades, all while communicating and sharing information in real time.

3.2 KEY COMPONENTS

The architecture of a multi-agent LLM system is built on several interconnected components, each playing a critical role in ensuring the system's functionality, efficiency, and adaptability. These components include agents, communication protocols, decision-making processes, knowledge bases, and interfaces. Together, they enable the system to operate collaboratively and intelligently in complex environments like traditional and decentralized finance.

3.2.1 Agents

Agents are the core building blocks of a multi-agent LLM system. Each agent is an autonomous entity powered by a large language model, equipped with specific roles, capabilities, and objectives. For example, in a financial context, one agent might specialize in analyzing market trends, another in assessing risk, and a third in executing trades. These agents can operate independently or collaboratively, depending on the task at hand. The strength of multi-agent systems lies in their ability to distribute tasks across multiple agents, enabling them to handle complex, multi-faceted problems more effectively than a single agent could. Each agent is designed to perform its role with a high degree of accuracy and efficiency, leveraging the advanced reasoning and language capabilities of LLMs to process information, make decisions, and communicate with other agents.

Let A=$a_1, a_2, \dots, a_n$ represent the set of agents in the multi-agent system, where each agent a_i is assigned a specific role R_i. The role $R_i$defines the agent's responsibilities and capabilities, such as data analysis, risk assessment, or trade execution.

$$R_i: a_i \rightarrow Task_i$$

where $Task_i$ represents the specific task or function performed by agent a_i.

3.2.2 Communication Protocols

Before jumping to decision making, we first check communication protocols. Communication protocols are the mechanisms through which agents interact and share information. These protocols define the rules and formats for exchanging data, ensuring that agents can understand and respond to each other effectively. Communication can occur through various methods, such as message passing, shared memory, or even natural language conversations. For instance, in a financial trading system, one agent might send a message to another agent requesting an analysis of market conditions, and the receiving agent might respond with a detailed report. The choice of communication protocol depends on the system's design and requirements. Effective communication is essential for coordination and collaboration among agents, enabling them to work together seamlessly and achieve shared goals.

Agents communicate through a set of protocols *P*, which define the rules and formats for exchanging information. Let M_{ij} represent the message sent from agent a_i to agent a_j.

$$M_{ij} = \left(a_i, a_j, Content_{ij}, Timestamp_{ij}\right)$$

where $Content_{ij}$ is the information exchanged (e.g., market data, risk metrics). $Timestamp_{ij}$ is the time at which the message is sent. The communication protocol *P* ensures that messages are transmitted and interpreted correctly:

$$P: M_{ij} \rightarrow ValidMessage$$

3.2.3 Decision-Making Processes

Decision-making processes in a multi-agent LLM system can be either centralized or decentralized, depending on the system's architecture and objectives. In a centralized system, a coordinator agent oversees the decision-making process, collecting inputs from other agents and making final decisions. This approach is useful in scenarios where a single point of control is necessary, such as regulatory compliance or risk management. In contrast, decentralized systems rely on consensus algorithms or negotiation mechanisms to enable agents to reach decisions collectively. This approach is more suitable for dynamic and distributed environments, such as decentralized finance, where no single entity has control over the entire system. Regardless of the approach, decision-making in multi-agent LLMs is typically data-driven, leveraging the vast amounts of information processed by the agents to make informed and accurate decisions.

Journal of Industrial Engineering and Applied Science
Journal Home: http://jieas.suaspress.org/ | CODEN: JIEAAE
Vol. 3, No. 1, 2025 | ISSN 3005-6071 (Print) | ISSN 3005-608X (Online)

The decision-making process D can be modeled as a function that takes input from multiple agents and produces an output decision $Decision_k$.

$$D: (Input_1, Input_2, \ldots, Input_n) \rightarrow Decision_k$$

where:

- $Input_i$ represents the input provided by agent a_i.
- $Decision_k$ is the final decision made by the system.

For decentralized decision-making, a consensus mechanism CC can be used:

$$C: (Input_1, Input_2, \ldots, Input_n) \rightarrow ConsensusDecision$$

3.2.4 Knowledge Base

The knowledge base is a centralized or distributed repository of information that agents can access and update. It serves as the system's memory, storing domain-specific data, historical records, and contextual information. In a financial context, the knowledge base might include market data, regulatory guidelines, transaction histories, and risk models. Agents query the knowledge base to retrieve relevant information and update it with new insights or findings. For example, an agent analyzing market trends might retrieve historical price data from the knowledge base, while an agent assessing risk might update the knowledge base with new risk metrics. The knowledge base ensures that all agents have access to the same information, promoting consistency and coherence in their actions and decisions.

3.2.5 Interfaces

Interfaces are the points of interaction between the multi-agent LLM system and external entities, such as humans or other systems. These interfaces enable seamless communication and integration, ensuring that the system can operate effectively within its environment. Human-agent interfaces, such as chatbots or dashboards, allow users to interact with the system, providing input, receiving outputs, and monitoring performance. For example, a portfolio manager might use a dashboard to view real-time updates on asset performance and adjust investment strategies accordingly. System integration interfaces, such as APIs and connectors, enable the system to interact with external platforms, such as trading systems, blockchain networks, or regulatory databases. These interfaces ensure that the multi-agent LLM system can function as part of a larger ecosystem, exchanging data and coordinating actions with other systems.

3.3 Applications in TradFi

Multi-agent large language models have the potential to revolutionize traditional finance TradFi by enhancing decision-making, automating complex processes, and improving operational efficiency. These systems can address some of the most pressing challenges in TradFi, such as inefficiencies in transaction processing, the need for real-time risk assessment, and the growing complexity of regulatory compliance. Below, we explore three key applications of multi-agent LLMs in TradFi: portfolio management, risk assessment and fraud detection, and regulatory compliance.

3.3.1 Portfolio Management

On both the buy-side and sell-side business, portfolio management is a critical function in TradFi, involving the selection and management of investment assets to achieve specific financial goals. Multi-agent LLMs can significantly enhance this process by enabling real-time analysis, optimization, and execution of investment strategies. In a multi-agent system, different agents can be assigned specialized roles, such as analyzing market trends, assessing risk, and executing trades. For example, one agent might monitor macroeconomic indicators and global market trends, while another evaluates the financial health and performance of individual companies. A third agent could then use this information to optimize asset allocation and execute trades in real time. By leveraging the collective intelligence of multiple agents, multi-agent LLMs can provide more accurate and timely insights, enabling portfolio managers to make better-informed decisions and achieve superior returns.

3.3.2 Risk Assessment and Fraud Detection

Risk assessment and fraud detection help institutions identify and mitigate potential threats to their operations. Multi-agent LLMs can enhance these processes by enabling collaborative analysis of vast amounts of data from diverse sources. For instance, one agent analyzes transaction data to identify unusual patterns or anomalies, while another assesses customer behavior to detect potential signs of fraud. A third agent could then correlate these findings with external data, such as market trends or regulatory alerts, to provide a comprehensive risk assessment. By working together, these agents can identify risks and fraudulent activities more quickly and accurately than traditional methods, reducing the likelihood of financial losses and reputational damage. Moreover, the adaptability of multi-agent LLMs allows them to continuously learn and improve, staying ahead of emerging threats and evolving fraud tactics.

3.3.3 Regulatory Compliance

Regulatory compliance is a major challenge for financial institutions, requiring them to adhere to a complex and ever-changing set of rules and guidelines. Multi-agent LLMs can streamline this process by automating the monitoring, analysis, and reporting of compliance-related data. For example, one agent might track regulatory changes and updates, ensuring that the institution is aware of new requirements. Another agent could audit transactions and activities to ensure compliance with these regulations, flagging any potential violations for further review. A third agent might then generate reports and documentation for regulators, providing a transparent and auditable record of the institution's compliance efforts. By automating these tasks, multi-agent LLMs can reduce the burden on compliance teams, minimize the risk of non-compliance, and ensure that institutions remain in good standing with regulators.

Additionally, the system's ability to process and analyze large volumes of data in real time enables it to identify and address compliance issues more quickly and effectively than traditional methods.

3.4 Applications in DeFi

Unlike TradFi, DeFi represents a paradigm shift in the financial industry, leveraging blockchain technology to create an open, permissionless, and transparent financial ecosystem. However, the rapid growth of DeFi has introduced unique challenges, such as smart contract vulnerabilities, liquidity fragmentation, and regulatory uncertainty. Multi-agent large language models offer solutions to these challenges by enabling collaborative, intelligent, and adaptive systems. Below, we explore three key applications of multi-agent LLMs in DeFi: smart contract optimization, decentralized autonomous organizations (DAOs), and cross-chain interoperability. We will illustrate those points one by one.

3.4.1 Smart Contract Optimization

Smart contracts are the backbone of DeFi, enabling the automation of financial services such as lending, borrowing, and trading. However, smart contracts are prone to vulnerabilities, which can lead to significant financial losses and security breaches. Multi-agent LLMs can enhance the security and efficiency of smart contracts by enabling collaborative auditing and optimization. For example, one agent might analyze the code of a smart contract to identify potential vulnerabilities, such as reentrancy attacks or integer overflows. Another agent could simulate the execution of the contract under various conditions to assess its robustness and performance. A third agent might then propose optimizations or patches to address any identified issues. By working together, these agents can ensure that smart contracts are secure, efficient, and reliable, reducing the risk of exploits and enhancing the overall stability of the DeFi ecosystem.

3.4.2 Decentralized Autonomous Organizations

Decentralized autonomous organizations are key innovations in DeFi, enabling decentralized governance and decision-making. However, managing DAO can be challenging, particularly as the number of participants and proposals grows. Multi-agent LLMs can facilitate the governance and decision-making processes in DAOs by providing intelligent analysis and automation. For instance, one agent might analyze proposals and assess their potential impact on the DAO's objectives. Another agent could evaluate voter sentiment and preferences, providing insights into the likely outcome of a vote. A third agent might then execute the decisions made by the DAO, ensuring that they are implemented accurately and efficiently. By leveraging the collective intelligence of multiple agents, multi-agent LLMs can enhance the transparency, efficiency, and effectiveness of DAO governance, enabling more informed and democratic decision-making.

3.4.3 Cross-Chain Interoperability

Cross-chain interoperability is a critical challenge in DeFi, as the ecosystem is fragmented across multiple blockchain networks, each with its own protocols and standards. Multi-agent LLMs can enable seamless transactions and interactions across different blockchain networks, enhancing the liquidity and functionality of DeFi. For example, one agent might monitor transactions and activities on the Ethereum network, while another does the same on the Binance Smart Chain. A third agent could then facilitate the transfer of assets between these networks, ensuring that transactions are carried out accurately and efficiently. By enabling cross-chain interoperability, multi-agent LLMs can create a more connected and cohesive DeFi ecosystem, allowing users to access a wider range of services and opportunities. This interoperability also enhances the scalability and resilience of DeFi, as it reduces the reliance on any single blockchain network and enables the system to adapt to changing conditions and demands.

4 CASE STUDIES AND USE CASES

4.1 Case Study 1: Multi-Agent LLMs in TradFi – Automated Trading Systems

Automated trading systems are widely used in traditional finance to execute trades based on predefined algorithms and market conditions. However, these systems often struggle to adapt to rapidly changing market dynamics and complex multi-variable scenarios. A multi-agent LLM system was implemented to enhance the performance of an automated trading system for a mid-sized hedge fund. The system consisted of three agents: one for market analysis, one for risk assessment, and one for trade execution.

Table 1: Performance Comparison of Automated Trading Systems

Metric	Traditional System	Multi-Agent LLM System	Improvement
Average Return (%)	8.5	12.3	+3.8
Risk-Adjusted Return (Sharpe Ratio)	1.2	1.8	+0.6
Trade Execution Speed (ms)	120	80	-40
Error Rate (%)	2.1	0.7	-1.4

Table 1 demonstrate the multi-agent LLM system having significant improvements over the traditional system. The average return increased from 8.5% to 12.3%, while the risk-adjusted return (measured by the Sharpe Ratio) improved from 1.2 to 1.8. The trade execution speed also decreased from 120 milliseconds to 80 milliseconds, reducing latency and improving efficiency. Additionally, the error rate dropped from 2.1% to 0.7%, highlighting the system's ability to make more accurate and informed decisions. These results underscore the advantages of multi-agent LLMs in handling traditional financial trading, like power-type derivatives for rough volatility with jumps introduced in [47].

4.2 Case Study 2: Multi-Agent LLMs in DeFi – Decentralized Lending Protocols

Decentralized lending protocols are a cornerstone of DeFi, enabling users to borrow and lend assets without intermediaries. However, these protocols face challenges such as liquidity fragmentation, interest rate volatility, and smart contract vulnerabilities. A multi-agent LLM system was deployed to optimize a decentralized lending protocol on the Ethereum blockchain. The system included agents for liquidity management, risk assessment, and smart contract auditing.

Table 2: Performance Comparison of Decentralized Lending Protocols

Metric	Traditional Protocol	Multi-Agent LLM Protocol	Improvement
Liquidity Utilization (%)	65	82	+17
Interest Rate Volatility (%)	12.4	8.2	-4.2
Smart Contract Vulnerabilities Detected	3	1	-2
User Satisfaction (Scale: 1-10)	6.8	8.5	+1.7

The multi-agent LLM protocol outperformed the traditional protocol across all key metrics (See Table 2). Liquidity utilization increased from 65% to 82%, indicating more efficient use of available assets. Interest rate volatility decreased from 12.4% to 8.2%, providing greater stability for borrowers and lenders. The system also detected and resolved two additional smart contract vulnerabilities, reducing the risk of exploits. User satisfaction, measured on a scale of 1 to 10, improved from 6.8 to 8.5, reflecting the system's ability to deliver a more reliable and user-friendly experience. Multi-agent LLMs can enhance the efficiency, security, and usability of decentralized lending protocols. In the future, we can even combine it with regulated contrastive learning defined in [48]. It also strike a balance between liquidity utilization and vulnerability detection.

4.3 Use Case 1: Cross-Border Payments

Cross-border payments are a critical component of global finance, but they are often plagued by high costs, slow processing times, and regulatory complexities. A multi-agent LLM system was implemented to streamline cross-border payments for a multinational corporation. The system included agents for regulatory compliance, currency conversion, and transaction routing.

Table 3 Performance Comparison of Cross-Border Payment Systems

Metric	Traditional System	Multi-Agent LLM System	Improvement
Transaction Cost (%)	5.2	3.1	-2.1
Processing Time (hours)	24	8	-16
Compliance Errors (%)	4.5	1.2	-3.3
Success Rate (%)	92	98	+6

Table 3 shows that the multi-agent LLM system significantly improved the efficiency and reliability of cross-border payments. Transaction costs decreased from 5.2% to 3.1%, while processing times were reduced from 24 hours to just 8 hours. Compliance errors dropped from 4.5% to 1.2%, ensuring that transactions adhered to regulatory requirements. The success rate of transactions also increased from 92% to 98%, highlighting the system's ability to handle complex, multi-jurisdictional payments with greater accuracy and reliability. Multi-agent LLMs addresses the challenges of cross-border payments, particularly in terms of cost, speed, and compliance.

4.4 Use Case 2: Regulatory Compliance

Regulatory compliance requires navigating a complex and ever-changing landscape of rules and guidelines. A multi-agent LLM system was deployed to automate compliance monitoring for a large bank. The system included agents for regulatory tracking, transaction auditing, and report generation. Most of these systems have two parts, a regulatory detection part and a compliance convergence part.

Table 4 Performance Comparison of Regulatory Compliance Systems

Metric	Traditional System	Multi-Agent LLM System	Improvement
Compliance Violations Detected (%)	85	95	+10
Time to Generate Reports (hours)	12	3	-9
False Positives	15	5	-10

Metric	Traditional System	Multi-Agent LLM System	Improvement
(%)			
Cost of Compliance (USD per transaction)	0.50	0.30	-0.20

I Table 4, the multi-agent LLM system detected 95% of compliance violations, compared to 85% for the traditional system, while reducing false positives from 15% to 5%. The time required to generate compliance reports decreased from 12 hours to just 3 hours, enabling faster and more efficient reporting. Additionally, the cost of compliance per transaction dropped from 0.50 to 0.30, highlighting the system's ability to reduce operational costs while maintaining high levels of accuracy and reliability. These results underscore the potential of multi-agent LLMs to transform regulatory compliance, making it more efficient, accurate, and cost-effective. It can detect violations that appear in both centralized and de-centralized form.

5 FUTURE DIRECTIONS AND CONCLUSION

One promising direction is the integration of multi-agent LLMs with emerging technologies such as quantum computing and the Internet of Things (IoT). Quantum computing could enhance the computational power of these systems, enabling them to solve complex financial problems more efficiently, while IoT devices could provide real-time data streams for more accurate decision-making. Another area of focus is improving the interpretability and transparency of multi-agent LLMs, which is critical for gaining the trust of stakeholders and ensuring ethical use. Researchers could explore techniques such as explainable AI to make the decision-making processes of these systems more understandable. The multi-agent LLM can possibly leverage on the web-based question answering systems introduced in [49].

From a practical standpoint, the adoption of multi-agent LLMs in finance will require collaboration among various stakeholders, including financial institutions, regulators, and technology providers. Financial institutions will need to invest in infrastructure and training to integrate these systems into their operations, while regulators will need to develop frameworks to ensure their safe and ethical use. Policymakers could play a key role in fostering innovation by creating sandboxes for testing multi-agent LLMs in controlled environments. For end-users, the widespread adoption of these systems could lead to more personalized and accessible financial services, reducing barriers to entry and promoting financial inclusion.

Through the proposed conceptual framework, we have shown how multi-agent LLMs can enhance decision-making, automation, and interoperability in financial ecosystems. Empirical results from case studies highlight significant improvements in key performance metrics: in portfolio management, multi-agent LLMs increased average returns by 3.8% and improved the Sharpe Ratio by 0.6; in DeFi, smart contract optimization reduced vulnerabilities by 66%, and decentralized lending protocols achieved a 17% increase in liquidity utilization. These findings underscore the ability of our multi-agent LLMs to deliver greater efficiency, accuracy, and scalability compared to traditional approaches.

ACKNOWLEDGMENTS

The authors thank the editor and anonymous reviewers for their helpful comments and valuable suggestions.

FUNDING

Not applicable.

INSTITUTIONAL REVIEW BOARD STATEMENT

Not applicable.

INFORMED CONSENT STATEMENT

Not applicable.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

CONFLICT OF INTEREST

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

PUBLISHER'S NOTE

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

AUTHOR CONTRIBUTIONS

Not applicable.

ABOUT THE AUTHORS

DE LA CRUZ, Antio

Bloomberg Bazil AI Lab, Brazil.

REFERENCES

[1] J. Begenau, M. Farboodi and L. Veldkamp, "Big data in finance and the growth of large firms," Journal of Monetary Economics, vol. 97, p. 71–87, 2018.

[2] Z. Li, B. Wang and Y. Chen, "Incorporating economic indicators and market sentiment effect into US Treasury bond yield prediction with machine learning," Journal of Infrastructure, Policy and Development, vol. 8, p. 7671, 2024.

[3] F. Schär, "Decentralized finance: On blockchain-and smart contract-based financial markets," FRB of St. Louis Review, 2021.

[4] L. Kiffer, D. Levin and A. Mislove, "Analyzing ethereum's contract topology," in Proceedings of the Internet Measurement Conference 2018, 2018.

[5] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell and others, "Language models are few-shot learners," Advances in neural information processing systems, vol. 33, p. 1877–1901, 2020.

[6] C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature machine intelligence, vol. 1, p. 206–215, 2019.

[7] K. Amaruchkul, "Dynamic Network for Air Freight Forwarder's Stochastic Capacity Management," Science & Technology Asia, p. 36–56, 2024.

[8] S. Diao, C. Wei, J. Wang and Y. Li, "Ventilator pressure prediction using recurrent neural network," arXiv preprint arXiv:2410.06552, 2024.

[9] X. Tang, Z. Wang, X. Cai, H. Su and C. Wei, "Research on heterogeneous computation resource allocation based on data-driven method," in 2024 6th International Conference on Data-driven Optimization of Complex Systems (DOCS), 2024.

[10] Z. Wang, Y. Chen, F. Wang and Q. Bao, "Improved Unet model for brain tumor image segmentation based on ASPP-coordinate attention mechanism," arXiv preprint arXiv:2409.08588, 2024.

[11] Z. Wu, "Mpgaan: Effective and efficient heterogeneous information network classification," Journal of Computer Science and Technology Studies, vol. 6, p. 08–16, 2024.

[12] W. Zhang, J. Huang, R. Wang, C. Wei, W. Huang and Y. Qiao, "Integration of Mamba and Transformer–MAT for Long-Short Range Time Series Forecasting with Application to Weather Dynamics," arXiv preprint arXiv:2409.08530, 2024.

[13] H. Cao, Z. Zhang, X. Li, C. Wu, H. Zhang and W. Zhang, "Mitigating Knowledge Conflicts in Language Model-Driven Question Answering," arXiv preprint arXiv:2411.11344, 2024.

[14] J. Zhang, W. Zhang, C. Tan, X. Li and Q. Sun, "YOLO-PPA based Efficient Traffic Sign Detection for Cruise Control in Autonomous Driving," arXiv preprint arXiv:2409.03320, 2024.

[15] H. Li and W. Zhang, "Cat face recognition using Siamese network," in International Conference on Artificial Intelligence and Intelligent Information Processing (AIIIP 2022), 2022.

[16] M. Shen, J. Y. Zhang, L. Chen, W. Yan, N. Jani, B. Sutton and O. Koyejo, "Labeling cost sensitive batch active learning for brain tumor segmentation," in 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), 2021.

[17] W. Yan, E. Wu, A. G. Schwing and E. Rosenbaum, "Semantic Autoencoder for Modeling BEOL and MOL Dielectric Lifetime Distributions," in 2023 IEEE International Reliability Physics Symposium (IRPS), 2023.

[18] H.-C. Dan, Z. Huang, B. Lu and M. Li, "Image-driven prediction system: Automatic extraction of aggregate gradation of pavement core samples integrating deep learning and interactive image processing framework," Construction and Building Materials, vol. 453, p. 139056, 2024.

[19] H.-C. Dan, B. Lu and M. Li, "Evaluation of asphalt pavement texture using multiview stereo reconstruction based on deep learning," Construction and Building Materials, vol. 412, p. 134837, 2024.

[20] H.-C. Dan, P. Yan, J. Tan, Y. Zhou and B. Lu, "Multiple distresses detection for Asphalt Pavement using improved you Only Look Once Algorithm based on convolutional neural network," International Journal of Pavement Engineering, vol. 25, p. 2308169, 2024.

[21] Q. Yu, Z. Ke, G. Xiong, Y. Cheng and X. Guo, "Identifying Money Laundering Risks in Digital Asset Transactions Based on AI Algorithms," 2025.

[22] Z. Ke, S. Zhou, Y. Zhou, C. H. Chang and R. Zhang, "Detection of AI Deepfake and Fraud in Online Payments Using GAN-Based Models," arXiv preprint arXiv:2501.07033, 2025.

[23] K. Li, J. Chen, D. Yu, T. Dajun, X. Qiu, L. Jieting, S. Baiwei, Z. Shengyuan, Z. Wan, R. Ji and others, "Deep reinforcement learning-based obstacle avoidance for robot movement in warehouse environments," arXiv preprint arXiv:2409.14972, 2024.

[24] K. Li, L. Liu, J. Chen, D. Yu, X. Zhou, M. Li, C. Wang and Z. Li, "Research on reinforcement learning based warehouse robot navigation algorithm in complex warehouse layout," arXiv preprint arXiv:2411.06128, 2024.

[25] K. Li, J. Wang, X. Wu, X. Peng, R. Chang, X. Deng, Y. Kang, Y. Yang, F. Ni and B. Hong, "Optimizing automated picking systems in warehouse robots using machine learning," arXiv preprint arXiv:2408.16633, 2024.

[26] M. Wang, H. Adel, L. Lange, J. Strötgen and H. Schütze, "NLNDE at SemEval-2023 Task 12: Adaptive Pretraining and Source Language Selection for Low-Resource Multilingual Sentiment Analysis," in Proceedings of the 17th International Workshop on Semantic Evaluation (SemEval-2023), 2023.

[27] M. Wang, H. Adel, L. Lange, J. Strötgen and H. Schütze, "Rehearsal-Free Modular and Compositional Continual Learning for Language Models," in Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers), 2024.

[28] J. Xing, C. Gao and J. Zhou, "Weighted fuzzy rough sets-based tri-training and its application to medical diagnosis," Applied Soft Computing, vol. 124, p. 109025, 2022.

[29] C. Gao, J. Zhou, J. Xing and X. Yue, "Parameterized maximum-entropy-based three-way approximate attribute reduction," International Journal of Approximate Reasoning, vol. 151, p. 85–100, 2022.

[30] S. Feng, X. Chen and S. Li, "Wavelet Guided Visual State Space Model and Patch Resampling Enhanced U-Shaped Structure for Skin Lesion Segmentation," IEEE Access, vol. 12, pp. 181521-181532, 2024.

[31] S. Feng, R. Song, S. Yang and D. Shi, "U-net Remote Sensing Image Segmentation Algorithm Based on Attention Mechanism Optimization," in 2024 9th International Symposium on Computer and Information Processing Technology (ISCIPT), 2024.

[32] T. Hu, W. Zhu and Y. Yan, "Artificial intelligence aspect of transportation analysis using large scale systems," in Proceedings of the 2023 6th Artificial Intelligence and Cloud Computing Conference, 2023.

[33] S. S. Mehrkian and H. Davari-Ardakani, "An integrated model of sentiment analysis and quantitative index data for predicting stock market trends: A case study of the Tehran stock exchange," Expert Systems with Applications, p. 126298, 2024.

[34] W. Zhu and T. Hu, "Twitter Sentiment analysis of covid vaccines," in 2021 5th International Conference on Artificial Intelligence and Virtual Reality (AIVR), 2021.

[35] D. Liu, "Contemporary Model Compression on Large Language Models Inference," arXiv preprint arXiv:2409.01990, 2024.

[36] D. Liu, "MT2ST: Adaptive Multi-Task to Single-Task Learning," arXiv preprint arXiv:2406.18038, 2024.

[37] D. Liu, R. Waleffe, M. Jiang and S. Venkataraman, "GraphSnapShot: Graph Machine Learning Acceleration with Fast Storage and Retrieval," arXiv preprint arXiv:2406.17918, 2024.

[38] W. Xu, J. Chen and J. Xiao, "A Hybrid Price Forecasting Model for the Stock Trading Market Based on AI Technique," Authorea Preprints, 2024.

[39] W. Xu, J. Xiao and J. Chen, "Leveraging Large Language Models to Enhance Personalized Recommendations in E-commerce," arXiv preprint arXiv:2410.12829, 2024.

[40] J. Chen, J. Xiao and W. Xu, "A Hybrid Stacking Method for Short-Term Price Forecasting in Electricity Trading Market," in 2024 8th International Conference on Information Technology, Information Systems and Electrical Engineering (ICITISEE), 2024.

[41] W. You, J. Chen, H. Xie and Y. Ren, "Which uncertainty measure better predicts gold prices? New evidence from a CNN-LSTM approach," The North American Journal of Economics and Finance, p. 102375, 2025.

[42] Y. Zhang, P. Ji, A. Wang, J. Mei, A. Kortylewski and A. L. Yuille, "3D-Aware Neural Body Fitting for Occlusion Robust 3D Human Pose Estimation," in IEEE/CVF International Conference on Computer Vision (ICCV), 2023.

[43] Z. Li, J. H. Bookbinder and S. Elhedhli, "Optimal shipment decisions for an airfreight forwarder: Formulation and solution methods," Transportation Research Part C: Emerging Technologies, vol. 21, p. 17–30, 2012.

[44] W. Zhu, "Optimizing distributed networking with big data scheduling and cloud computing," in International Conference on Cloud Computing, Internet of Things, and Computer Applications (CICA 2022), 2022.

[45] J. Jin, H. Xu, P. Ji and B. Leng, "IMC-NET: Learning Implicit Field with Corner Attention Network for 3D Shape Reconstruction," in IEEE International Conference on Image Processing (ICIP), 2022.

[46] Y. Wang, Y. Zhang, M. Huo, R. Tian, X. Zhang, Y. Xie, C. Xu, P. Ji, W. Zhan, M. Ding and M. Tomizuka, "Sparse Diffusion Policy: A Sparse, Reusable, and Flexible Policy for Robot Learning," CoRR, vol. abs/2407.01531, 2024.

[47] L. Wang and W. Xia, "Power-type derivatives for rough volatility with jumps," Journal of Futures Markets, vol. 42, p. 1369–1406, 2022.

[48] Z. Li, B. Wang and Y. Chen, "A Contrastive Deep Learning Approach to Cryptocurrency Portfolio with US

Treasuries," Journal of Computer Technology and Applied Mathematics, vol. 1, p. 1–10, 2024.

[49] L. Wang, I. Lauriola and A. Moschitti, "Accurate training of web-based question answering systems with feedback from ranked users," in Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track), pp. 660-667,2023.