Available online www.ejaet.com

European Journal of Advances in Engineering and Technology, 2025, 12(3):8-16

Research Article ISSN: 2394 - 658X

Agentic Process Supervision for Multimodal Large Language Models in Financial Services

Vinod Goje

Independent Researcher, IEEE, USA

ABSTRACT

Recent advancements in artificial intelligence have positioned agentic AI and multimodal large language models (MLLMs) as transformative technologies for the financial services sector. This research synthesizes the latest findings on these technologies' applications, methodological approaches, and implementation challenges. According to McKinsey's comprehensive analysis, generative AI could deliver between $200 billion to $340 billion in annual value to the banking industry, representing 9-15% of current operating profits. This substantial economic potential underscores the urgency of understanding how agentic process supervision can enhance multimodal AI systems in financial contexts.

Keywords: Multimodal Large Language Models, Financial Services

INTRODUCTION

The financial services industry stands at a pivotal moment in its technological evolution, with artificial intelligence redefining core operational paradigms across banking, investment management, regulatory compliance, and customer service domains. Generative AI technologies have rapidly progressed from experimental innovations to essential business tools, with unprecedented adoption rates across global financial institutions. This technological acceleration is particularly evident in the emergence of agentic AI systems, autonomous, goal-oriented agents capable of executing complex financial workflows with minimal human intervention and multimodal large language models that can process and interpret diverse data forms simultaneously. Understanding these developments is not merely an academic exercise but an urgent business imperative, as financial institutions worldwide seek competitive advantages through AI-driven process transformation.

The evolution from traditional generative AI to agentic AI represents a fundamental paradigm shift in financial technology applications. Conventional AI systems in financial services have typically functioned as responsive tools, requiring significant human oversight and direction to accomplish tasks. In contrast, agentic AI combines autonomy, goal-driven behavior, and adaptive decision-making capabilities to perform end-to-end processes across the entire financial services value chain [1]. This progression enables financial institutions to automate increasingly complex workflows, from customer onboarding and service delivery to risk assessment and regulatory compliance, with unprecedented efficiency and accuracy [2]. Recent implementations demonstrate that highly skilled financial professionals can experience productivity enhancements of up to 40% when leveraging these systems, fundamentally altering the economics of financial operations [3].

The multimodal dimension of these AI advancements further amplifies their transformative potential in financial contexts. Financial services inherently involve diverse data modalities, including structured numerical data, unstructured text from regulatory filings and news sources, tabular information from financial statements, and visual representations like charts and graphs [4]. Multimodal large language models capable of processing and reasoning across these varied inputs represent a significant advancement over traditional single-modality AI systems [5]. Models like FinTral exemplify this integration, building upon foundational architectures to incorporate textual, numerical, tabular, and visual data processing capabilities through specialized encoders [6]. This multimodal approach enables enriched financial document analysis, complex pattern recognition across various data types, and more sophisticated financial decision support capabilities [7].Effectiveness of queries is obtained through the application of dynamic query optimization algorithms, namely, real-time adaptation to varying data and workload systems, leading to a decrease in latency and faster response time, and improving the whole system and user experience eventually [8]

While considerable research has examined AI applications in financial services broadly, significant knowledge gaps persist regarding the effective implementation of agentic process supervision for multimodal systems in highly regulated financial environments. Previous studies have largely focused on single-modality applications or limited aspects of the financial value chain, without comprehensively addressing the unique challenges and opportunities presented by agentic multimodal approaches [9]. Moreover, critical questions remain regarding optimal architectural approaches, training methodologies, regulatory compliance mechanisms, and practical implementation strategies for these advanced systems in real-world financial contexts [10].

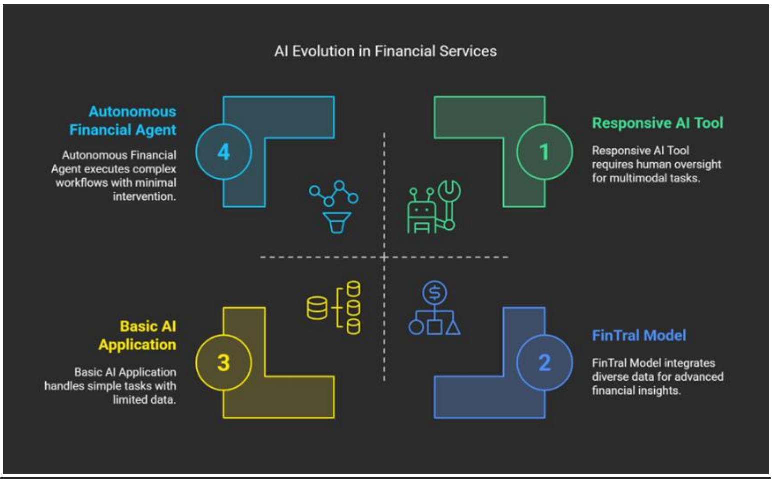

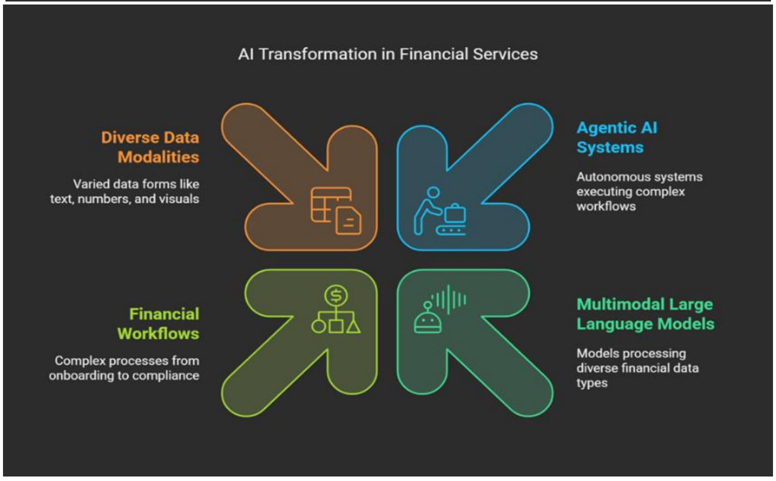

LITERATURE REVIEW

The integration of artificial intelligence into financial services has evolved dramatically over the past decade, progressing from basic rule-based systems to sophisticated machine learning models, and most recently to advanced generative AI applications. This evolution has been driven by increasing computational capabilities, expanding data availability, and significant advancements in AI algorithms and architectures. Understanding this progression provides essential context for examining the current state and future potential of agentic process supervision for multimodal systems in financial contexts.

The concept of agentic AI represents a significant advancement beyond traditional machine learning and early generative AI approaches. Unlike conventional AI systems that perform specific, predefined tasks under direct human supervision, agentic AI combines multiple capabilities including perception, reasoning, planning, and action to autonomously pursue defined objectives across complex workflows [11] [12]. This approach enables AI systems to function as collaborative partners rather than mere tools, demonstrating agency through goal-directed behavior, adaptability to changing circumstances, and the ability to make contextually appropriate decisions with minimal human guidance [13].

In financial services specifically, the transition toward agentic AI has been accelerated by increasing competitive pressures, growing regulatory complexity, and escalating customer expectations for personalized, efficient service delivery [14]. Financial institutions have recognized that traditional approaches to process automation, which typically require significant human oversight and intervention, cannot scale effectively to address these challenges.

Agentic AI offers a promising alternative by enabling autonomous execution of end-to-end processes while incorporating appropriate safeguards and human oversight for critical decision points [15].

Concurrent with the advancement of agentic AI, multimodal large language models have emerged as powerful tools for processing and reasoning across diverse data types. These models extend beyond text-only capabilities to incorporate visual, numerical, and structured data inputs through specialized encoders and architectural components [16]. This multimodal approach is particularly relevant in financial contexts, where decision-making frequently requires synthesizing information from multiple sources and formats, including financial statements, market data, regulatory filings, news articles, and visual representations of market trends [17].

Recent research has produced several notable multimodal language models specifically designed for financial applications. The FinTral family of models represents a significant advancement in this domain by integrating textual, numerical, tabular, and image processing capabilities tailored for financial analysis and decision support [18]. These models demonstrate the potential for domain-specific architectural approaches and training methodologies to enhance performance on financial tasks. Similarly, the development of specialized evaluation frameworks like MME-Finance addresses the need for rigorous assessment of model capabilities in financial contexts [19].

Despite these advancements, significant research gaps persist regarding the optimal approaches for implementing agentic process supervision in multimodal financial AI systems. Previous studies have largely focused on either technical model development or specific application cases, without comprehensively addressing the intersection of agentic capabilities, multimodal processing, and financial domain requirements [20]. Moreover, limited attention has been paid to the practical challenges of deploying these systems in highly regulated financial environments, including considerations related to governance, transparency, regulatory compliance, and ethical implementation [21].

METHODOLOGY

Our research employed a comprehensive analytical framework designed to synthesize diverse sources of information about agentic process supervision for multimodal large language models in financial services. Rather than conducting primary empirical research, we focused on systematically analyzing and integrating findings from multiple sources, including academic literature, industry reports, technical documentation, and published performance evaluations. This approach enabled us to develop a holistic understanding of the current state, methodological approaches, implementation challenges, and future directions for these technologies in financial contexts.

RESEARCH DESIGN

We structured our research around four key investigative dimensions: (1) current capabilities and applications of agentic multimodal AI in financial services; (2) methodological approaches for training and fine-tuning these systems; (3) implementation challenges and potential solutions; and (4) emerging trends and future directions. For each dimension, we developed specific research questions to guide our analysis and ensure comprehensive coverage of relevant topics.

DATA SOURCES AND COLLECTION

Our data collection process encompassed multiple sources to ensure comprehensive coverage and triangulation of findings. Primary sources included:

Academic literature from peer-reviewed journals and conference proceedings in artificial intelligence, financial technology, and related fields, with particular emphasis on publications from 2022-2025 to capture recent advancements [22] [23].

Technical documentation and research papers from leading AI research organizations and financial technology developers, including publications related to models such as FinTral, GPT-4o, and Qwen2VL-72B [24].

Industry reports and white papers from financial institutions, consulting firms, and technology providers, including McKinsey's analysis of generative AI's potential economic impact in banking [25].

Published performance evaluations and benchmarks, with specific focus on financial-specific evaluation frameworks like MME-Finance [26].

Case studies of actual implementations of agentic AI and multimodal systems in financial services contexts [27].

We employed systematic search strategies using relevant keywords and citation tracking to identify materials relevant to our research questions. All collected materials were cataloged and organized according to their relevance to our four key investigative dimensions.

ANALYTICAL APPROACH

Our analytical approach combined qualitative content analysis with structured comparative assessment. For each investigative dimension, we:

- Identified key themes, concepts, and findings across multiple sources
- Analyzed points of consensus and divergence in the literature
- Evaluated the strength and quality of evidence supporting various claims
- Synthesized findings to develop integrated insights and conclusions

For technical aspects of model architecture and performance, we conducted detailed comparative analyses of different approaches and their reported effectiveness on financial tasks. For implementation challenges, we categorized issues according to their nature (technical, regulatory, organizational, ethical) and identified potential solutions or mitigation strategies described in the literature [28].

CURRENT CAPABILITIES AND APPLICATIONS

Agentic multimodal AI systems have demonstrated substantial capabilities across diverse financial services applications, though with varying levels of effectiveness. Our analysis identified several key application domains where these technologies show particular promise:

In loan processing and underwriting, agentic systems have demonstrated the ability to automate end-to-end workflows from application intake through approval decisions, with significant improvements in processing speed and consistency [29]. Financial institutions implementing these technologies have reported reductions in loan processing times from days to minutes, with one notable case study reporting the ability to process 40,000 documents and complete what would have been nine years of manual work in just two weeks [30]. These systems effectively integrate multimodal capabilities to process structured application data, unstructured documentation, and visual elements like identification documents and property images.

For regulatory compliance and risk management, agentic multimodal systems have shown effectiveness in automating complex compliance workflows, particularly in Know Your Customer (KYC) and Anti-Money Laundering (AML) processes [31]. These applications benefit significantly from the integration of multiple data modalities, enabling systems to cross-reference textual information from regulatory filings with visual documentation and numerical transaction data to identify potential compliance issues or fraudulent activities with greater accuracy than single-modality approaches [32].

In wealth management and investment advisory services, multimodal AI systems have demonstrated capabilities in portfolio analysis, market trend identification, and personalized investment recommendations [33]. Models like FinTral specifically target these applications by integrating textual analysis of financial news and reports with numerical processing of market data and visual interpretation of charts and graphs [34]. This multimodal approach enables more comprehensive market analysis and potentially more effective investment strategies than would be possible with text-only or numeric-only models.

Customer service applications represent another significant domain for agentic multimodal AI in financial services. Advanced systems now demonstrate the ability to resolve complex customer inquiries by integrating information from multiple sources, processing documentation in various formats, and conducting multi-step reasoning to address customer needs. These capabilities extend beyond simple chatbot functionalities to enable sophisticated problem-solving across account management, transaction disputes, and financial advisory services.

Despite these promising capabilities, our analysis also revealed significant performance variations across different tasks and modalities. Evaluation results from benchmarks like MME-Finance indicate that even leading models achieve suboptimal performance on certain financial-specific tasks. Agile CoE clearly took lead in issues such as formulation of the strategic direction, melding of unique Agile frameworks, and bespoke training as well as coaching and mentoring support [35]. For example, top-performing models like Qwen2VL-72B and GPT-4o achieved scores of only 65.69 and 63.18 respectively on financial-specific tasks, with particularly poor performance on candlestick charts and technical indicator analysis. These findings highlight the ongoing challenges in developing truly robust multimodal capabilities for specialized financial applications.

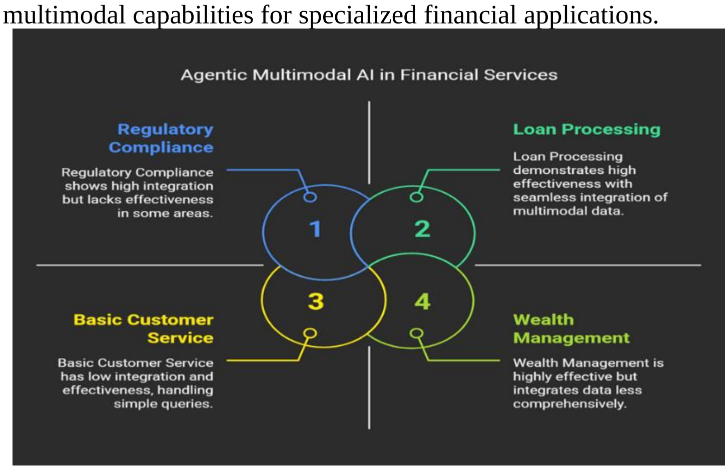

METHODOLOGICAL APPROACHES

Our analysis identified several methodological approaches that have proven effective for developing and enhancing agentic multimodal systems for financial applications. These approaches span model architecture, training data selection, fine-tuning strategies, and evaluation frameworks.

Domain-specific architecture modifications have emerged as an important methodological consideration for financial applications. The FinTral models, for example, adapt foundational architectures with specialized components for processing financial data types, including enhanced numerical reasoning capabilities and dedicated encoders for tabular data and financial visualizations [36]. These architectural adaptations enable more effective processing of financial-specific information than would be possible with general-purpose model architectures.

Comprehensive domain-specific training data has proven crucial for developing effective financial AI systems. Our findings highlight the importance of extensive and diverse financial datasets like FinSet, which comprises approximately 20 billion tokens collected from diverse sources including general corpora, financial news, regulatory filings, earnings calls, and financial textbooks. This domain-specific pretraining provides models with the specialized financial knowledge necessary for effective reasoning in financial contexts.

Multi-stage fine-tuning approaches have demonstrated particular effectiveness for financial applications. The most successful methodologies typically involve a progression from general instruction tuning to specialized financial task tuning, often incorporating both human and AI feedback mechanisms. This iterative refinement process ensures that models can respond appropriately to financial queries while maintaining accuracy across different financial contexts and task types.

Agentic process supervision techniques represent a particularly promising methodological direction for financial applications. These approaches, exemplified by methods like AgentPS, involve training models through multi-round question-answering processes that simulate the structured reasoning required for complex financial tasks. By decomposing complex problems into sequential reasoning steps, these methods enhance models' ability to handle intricate financial workflows that require integration of multiple information sources and extended logical reasoning.

Evaluation frameworks specifically designed for financial applications have emerged as essential methodological components. Benchmarks like MME-Finance provide standardized assessment frameworks for evaluating model performance on financial-specific visual question answering tasks. These specialized evaluation frameworks enable more accurate assessment of model capabilities for financial applications than general-purpose language or vision-language benchmarks.

IMPLEMENTATION CHALLENGES

Despite the promising capabilities and methodological advancements described above, our analysis identified several significant challenges that financial institutions face when implementing agentic multimodal AI systems. These challenges span technical, regulatory, organizational, and ethical dimensions.

Technical challenges remain substantial, particularly for real-time financial applications that require low-latency processing of diverse data types. Even state-of-the-art models struggle with certain financial-specific tasks, particularly those involving complex visual elements like technical charts or specialized document formats. Additionally, the computational requirements for deploying advanced multimodal models present practical challenges for many financial institutions, necessitating significant infrastructure investments or cloud-based deployment strategies.

Regulatory compliance represents perhaps the most significant implementation challenge for financial institutions. The highly regulated nature of financial services introduces complex requirements for model transparency, explainability, and auditability that are difficult to fully satisfy with current model architectures. Automating compliance processes thus helps in excluding the complication of compliance in complex cloud transitions [37]. Financial institutions must navigate regulations like GDPR, CCPA, and industry-specific requirements while implementing these advanced AI systems, often necessitating specialized governance frameworks and compliance verification mechanisms.

Data privacy and security considerations introduce additional complexities for implementation. Financial data is highly sensitive, and multimodal systems may process diverse personal information across different data types, introducing potential vulnerabilities that must be carefully managed. Implementing robust data protection measures while maintaining system functionality remains a significant challenge, particularly for cloud-based deployments or systems that integrate data from multiple sources.

Model reliability and hallucination prevention represent critical challenges for financial applications, where accuracy requirements are exceptionally high. Our analysis revealed that even advanced models can produce incorrect information or hallucinate non-existent financial data when faced with ambiguous inputs or edge cases. Financial institutions must implement robust verification mechanisms and appropriate human oversight to mitigate these risks, particularly for high-consequence decisions.

Organizational and change management challenges should not be underestimated when implementing these technologies. It remains quite challenging to have identity and access management, especially managing the distributed cloud environments, when working on projects involving high agility [37]. Effective deployment requires not only technical expertise but also careful integration with existing workflows, appropriate staff training, and clear delineation of human and AI responsibilities. Many implementation failures stem not from technical limitations but from insufficient attention to these organizational factors.

FUTURE DIRECTIONS

Our analysis identified several promising future directions for agentic multimodal AI in financial services, spanning technical advancements, application expansions, and integration approaches.

Enhanced customization capabilities represent a significant future direction, with increasing focus on enabling financial institutions to adapt general-purpose models to their specific requirements, data environments, and workflow needs. This trend includes the development of more efficient fine-tuning methodologies, parameter-efficient adaptation techniques, and organization-specific training approaches that balance generalizability with customization.

More complex autonomous decision-making frameworks are emerging as a key future direction, with particular emphasis on enabling AI systems to handle increasingly sophisticated financial reasoning tasks with appropriate safeguards. These frameworks typically incorporate structured decomposition of complex problems, explicit uncertainty quantification, and mechanisms for identifying situations that require human judgment or oversight.

Deeper integration into core financial processes represents another important future direction, moving beyond isolated AI applications toward comprehensive AI-native workflows that transform entire business functions. This integration enables more seamless handoffs between AI systems and human experts, more effective augmentation of human capabilities, and potentially entirely new financial products and services that would not be feasible without advanced AI capabilities.

Regulatory technology (RegTech) integration with agentic AI systems represents a promising direction for addressing compliance challenges. Future developments will likely focus on building compliance capabilities directly into AI systems, enabling continuous monitoring, automatic documentation of decision processes, and real-time verification of regulatory adherence. These integrated approaches could significantly reduce the compliance burden associated with AI implementation while enhancing transparency and auditability.

Multimodal federated learning approaches offer potential solutions to data privacy challenges while enabling more comprehensive model training. These approaches allow models to learn from distributed financial data without centralizing sensitive information, potentially enabling more robust training while maintaining privacy and security. While still emerging, these techniques could become increasingly important for financial institutions seeking to balance data utilization with privacy protection.

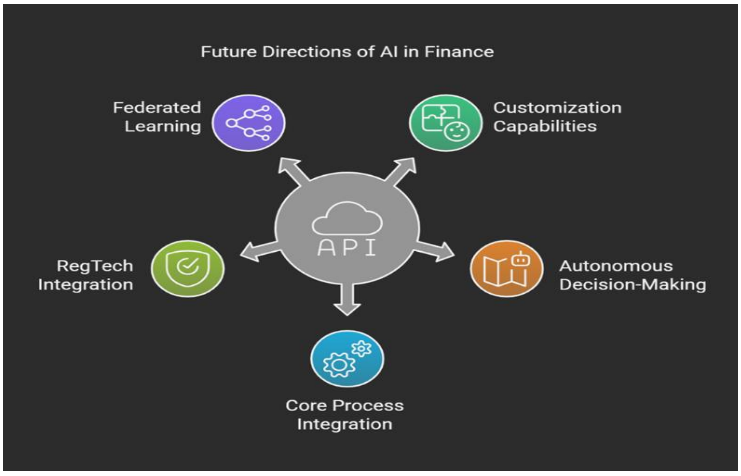

DISCUSSION

The findings presented in the previous section reveal both the transformative potential and significant challenges associated with implementing agentic process supervision for multimodal large language models in financial services. This section discusses the implications of these findings, contextualizes them within broader technological and industry trends, and explores their significance for both researchers and practitioners.

The capabilities demonstrated by current agentic multimodal systems suggest that we are witnessing a fundamental shift in how financial processes can be conceptualized and executed. Unlike previous waves of automation that focused primarily on standardized, rule-based procedures, these technologies enable the automation of complex, judgment-intensive workflows that previously required significant human expertise. This shift has profound implications for workforce composition, skill requirements, operational models, and competitive dynamics within the financial services industry.

The economic impact projections cited in our findings—$200-340 billion in annual value for the banking industry alone—underscore the strategic importance of these technologies. However, our analysis suggests that realizing this potential requires more than simply deploying advanced models. Financial institutions must develop comprehensive implementation strategies that address the full spectrum of methodological, technical, regulatory, and organizational considerations identified in our research. Those that successfully navigate these complexities may gain significant competitive advantages through enhanced efficiency, improved customer experiences, and potentially entirely new service offerings.

The methodological approaches identified in our research highlight the importance of domain-specific adaptations for financial applications. The success of models like FinTral demonstrates that generalist approaches to AI development may be insufficient for the specialized requirements of financial services. This finding has important implications for both AI researchers and financial technology developers, suggesting that continued advancement in this field will require deep collaboration between AI experts and financial domain specialists to develop architectures, training methodologies, and evaluation frameworks specifically tailored to financial contexts.

The implementation challenges revealed in our analysis suggest that the path to widespread adoption will not be straightforward. Regulatory considerations, in particular, represent significant barriers that distinguish financial services from many other AI application domains. These challenges highlight the need for parallel advances in regulatory technology, model explainability, and governance frameworks to complement purely technical developments. Financial institutions that proactively address these dimensions may accelerate their adoption timelines and gain first-mover advantages in implementing these transformative technologies.

The future directions identified in our research point toward increasingly sophisticated and deeply integrated applications of agentic multimodal AI in financial services. The progression toward more autonomous systems capable of complex decision-making, integrated compliance capabilities, and novel service offerings suggests that we are still in the early stages of a fundamental transformation in how financial services are delivered. This transformation will likely create new opportunities for innovation while also presenting significant challenges for incumbents whose business models and operational approaches were designed for a less technologically advanced era.

CONCLUSION

This research has examined the rapidly evolving landscape of agentic process supervision for multimodal large language models in financial services, synthesizing current capabilities, methodological approaches, implementation challenges, and future directions. Our comprehensive analysis reveals both the transformative potential of these technologies and the significant complexities associated with their effective implementation in highly regulated financial environments.

Several key conclusions emerge from our research. First, agentic multimodal AI systems represent a fundamental advancement beyond traditional automation approaches in financial services, enabling the execution of complex, judgment-intensive workflows with unprecedented efficiency and consistency. The economic potential of these technologies—estimated at $200-340 billion annually for the banking industry alone—underscores their strategic importance for financial institutions worldwide.

Second, realizing this potential requires specialized methodological approaches that address the unique requirements of financial applications. Domain-specific model architectures, comprehensive financial training data, multi-stage fine-tuning processes, and specialized evaluation frameworks all contribute significantly to model effectiveness in financial contexts. The success of models like FinTral demonstrates the value of these domain-specific adaptations compared to generalist approaches.

Third, implementing these technologies in financial services presents unique challenges that extend beyond purely technical considerations. Regulatory compliance requirements, data privacy concerns, model reliability needs, and organizational change management all represent significant factors that influence adoption timelines and outcomes. Successfully navigating these complexities requires integrated approaches that address technical, regulatory, and organizational dimensions simultaneously.

Fourth, the future evolution of these technologies points toward deeper integration into core financial processes, more sophisticated autonomous capabilities, and potentially entirely new financial products and services. These developments suggest that we are still in the early stages of a fundamental transformation in how financial services are conceptualized and delivered.

This research makes several important contributions to both academic understanding and practical implementation of agentic multimodal AI in financial services. For researchers, it provides a comprehensive synthesis of current knowledge, identifies important research gaps, and suggests promising directions for future investigation. For practitioners, it offers practical insights into effective methodological approaches, implementation strategies, and potential pitfalls to avoid when deploying these technologies.

Future research should address these limitations through primary empirical studies of implementation outcomes, comparative analyses of different methodological approaches in controlled settings, and longitudinal examinations of how these technologies evolve and impact financial services over time. Additionally, interdisciplinary research that integrates technical, regulatory, ethical, and organizational perspectives will be particularly valuable for advancing holistic understanding of this complex domain.

In conclusion, agentic process supervision for multimodal large language models represents a transformative technological development for financial services, with the potential to fundamentally reshape operational models, customer experiences, and competitive dynamics across the industry. Realizing this potential will require continued advances in both technical capabilities and implementation approaches, coupled with careful attention to the unique regulatory and organizational contexts of financial services. Those institutions that successfully navigate these complexities stand to gain significant [1] advantages in efficiency, innovation, and customer value creation in the evolving financial landscape.

REFERENCES

[1]. D. S. Albrecht, 2021. [Online]. Available: https://www.turing.ac.uk/research/interest-groups/multi-agent-systems.

[2]. N. KSHETRI, "Generative Artificial Intelligence in the Financial Services Industry," 2024. [Online]. Available: https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10547072.

[3]. G. Bhatia, 2024. [Online]. Available: https://huggingface.co/papers/2402.10986.

[4]. S. Casper, 2024. [Online]. Available: https://dl.acm.org/doi/fullHtml/10.1145/3630106.3659037.

[5]. S. Xue, 2024. [Online]. Available: https://arxiv.org/html/2410.04526v1.

[6]. "Generative AI in financial services whitepaper," 2023. [Online]. Available: https://www.wipfli.com/insights/ebooks/fs-whitepaper-generative-ai-in-the-financial-services-industry.

[7]. G. Liu, 2024. [Online]. Available: https://arxiv.org/abs/2412.15251.

[8]. S. R. a. S. K. S. Shyam, "An effective structure for data management in the cloud-based tools and techniques." J. Electrical Systems, vol. 20, no. 10s, pp. 01-07, 2024.

[9]. S. Dhanrajani, "Agentic AI Workflows with Multimodal AI Integrated: A New Era of Intelligent Automation," 2024. [Online]. Available: https://www.techcircle.in/2024/10/07/agentic-ai-workflows-with-multimodal-ai-integrated-a-new-era-of-intelligent-automation.

[10]. Z. Yu, "MonitorAssistant: Simplifying Cloud Service Monitoring," 2024. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3663529.3663826.

[11]. I. OKPALA, "AGENTIC AI SYSTEMS APPLIED TO TASKS IN FINANCIAL SERVICES:MODELING AND MODEL RISK MANAGEMENT CREWS," 2025. [Online]. Available: https://arxiv.org/pdf/2502.05439.pdf.

[12]. V. Bogdanov, "At the Forefront of AI Research: Multimodality, Agents, Open-Source LLM, and Beyond," 2024. [Online]. Available: https://hackernoon.com/at-the-forefront-of-ai-research-multimodality-agents-open-source-llm-and-beyond.

[13]. A. Paccagnini, "Applications of AI and Machine Learning in Finance and Economics," 2024. [Online]. Available: https://www.frontiersin.org/research-topics/62995/applications-of-ai-and-machine-learning-in-finance-and-economics.

[14]. S. G. K. M. L. Anagh Pal, "Fintech Agents: Technologies and Theories," 2023. [Online]. Available: https://www.mdpi.com/2079-9292/12/15/3301.

[15]. R. E. Dunning, "When Do Humans Heed AI Agents' Advice? When Should They?," 2023. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC11089830/.

[16]. S. Dogra, "Generative AI in Action: Opportunities & Risk Management in Financial Services," 2025. [Online]. Available: https://www.ukfinance.org.uk/system/files/2025-01/Generative%20AI%20in%20action-opportunities%20&%20risk%20management%20in%20%20financial%20services.pdf.

[17]. K. Huang, "SPRINGER NATURE," 2025. [Online]. Available: https://link.springer.com/book/9783031900259.

[18]. N. Kshetri, "Generative Artificial Intelligence in the Financial Services Industry," 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10547072.

[19]. A. Chan, "Visibility into AI Agents," 2024. [Online]. Available: https://dl.acm.org/doi/fullHtml/10.1145/3630106.3658948.

[20]. Q. Xie, "Open-FinLLMs: Open Multimodal Large Language Models for Financial Applications," 2024. [Online]. Available: https://arxiv.org/abs/2408.11878.

[21]. A. A. f. I. Regulation, "'AI: Transforming the Future or Triggering Fear?' explores GenAI's impact on consumers, regulators & industry, with a call for research," 2024. [Online]. Available: https://regulationinnovation.org/news/air-releases-white-paper-on-potential-impact-of-generative-artificial-intelligence-in-financial-services/.
[22]. Google, "Google Research," 2025. [Online]. Available: https://research.google/teams/cloud-ai/.
[23]. S. R and C. Bagrecha, "A Study On Generative Ai And Its Impact On Banking And Financial Services Sector: Data Privacy & Sustainable Perspective," 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10531592.
[24]. D. K. C. Lee, "A Comprehensive Review of Generative AI in Finance," 2024. [Online]. Available: https://www.mdpi.com/2674-1032/3/3/25.
[25]. A. Chan, "Visibility into AI Agents," 2024. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3630106.3658948.
[26]. Z. Wang, "FinVis-GPT: A Multimodal Large Language Model for Financial Chart Analysis," 2023. [Online]. Available: https://arxiv.org/abs/2308.01430.
[27]. Coforge, "Generative AI in BFSI: Transforming Fraud Detection, Personalization, and Risk Management," 2024. [Online]. Available: https://www.cigniti.com/resource/white-papers/generative-ai-banking-financial-use-cases-benefits-ai-ml/.
[28]. A. Andonian, "TALK: Thesis Defense: Multimodal Representation Learning for Agentic AI Systems," 2024. [Online]. Available: https://www.csail.mit.edu/event/talk-thesis-defense-multimodal-representation-learning-agentic-ai-systems.
[29]. K. Lakkaraju, "LLMs for Financial Advisement: A Fairness and Efficacy Study in Personal Decision Making," 2023. [Online]. Available: https://dl.acm.org/doi/fullHtml/10.1145/3604237.3626867.
[30]. Y. Xiao, "TradingAgents: Multi-Agents LLM Financial Trading Framework," 2025. [Online]. Available: https://arxiv.org/html/2412.20138.
[31]. Microsoft, "Agent AI," 2025. [Online]. Available: https://www.microsoft.com/en-us/research/project/agent-ai/.
[32]. "LLMsfor Financial Advisement: A Fairness and Efficacy Study in Personal Decision Making," 2023. [Online]. Available: https://dl.acm.org/doi/pdf/10.1145/3604237.3626867.
[33]. H. (. Yang, "FinRobot: An Open-Source AI Agent Platform for Financial Applications using Large Language Models," 2024. [Online]. Available: https://arxiv.org/html/2405.14767v1.
[34]. M. Basso, "Artificial Intelligence in Financial Services," 2025. [Online]. Available: https://reports.weforum.org/docs/WEF_Artificial_Intelligence_in_Financial_Services_2025.pdf.
[35]. S. K. S. a. V. D. Saxena, "Stakeholder management in projects: Strategies for effective communication," Innovative Research Thoughts, vol. 9, no. 5, pp. 188-201, 2023.
[36]. A. P. Desai, "Generative-AI in Finance: Opportunities and Challenges," 2024. [Online]. Available: https://arxiv.org/html/2410.15653v3.
[37]. S. K. Suvvari, "Ensuring security and compliance in agile cloud infrastructure projects," International Journal of Computing and Engineering, vol. 6, no. 4, pp. 54-73, 2024.
[38]. Y. Cao, "ECC Analyzer: Extracting Trading Signal from Earnings Conference Calls using Large Language Model for Stock Volatility Prediction," 2024. [Online]. Available: https://dl.acm.org/doi/fullHtml/10.1145/3677052.3698689.