

Survey of Agents Methodologies in the Financial Domain

Mohammad Luqman
Independent Researcher
Unaffiliated
luqmankgp@gmail.com

Himaanshu Gauba
Independent Researcher
Unaffiliated
gauba.himanshu@gmail.com

Prabhat Kumar
Independent Researcher
Unaffiliated
dawnavd@gmail.com

Akshar Prabhu Desai
Independent Researcher
Unaffiliated
akshar.iitb@gmail.com

Ajay Yadav
Independent Researcher
Unaffiliated
ajayyadav154@gmail.com

Ritu Prajapati
Independent Researcher
Unaffiliated
rituprajapati53@gmail.com

Pranjul Yadav
Independent Researcher
Unaffiliated
ipranjul@gmail.com

Abstract—Agentic AI techniques are autonomous systems that can make their decisions independently to work towards a specific goal in a dynamic environment. This paper explores the emerging field of Agentic AI in finance, examining the opportunities it presents through automation, faster decision making, and personalization along with the challenges. In addition, we also discuss the various techniques used to develop such agents along with their inherent design patterns, composite structures, and frameworks used to develop and evaluate such agents. Lastly, we discuss the future and emerging directions within this research area by focusing on multiagent and human-agent collaboration. We believe that this manuscript will provide a comprehensive overview of work done in the cross-disciplinary area of Agentic AI and finance.

Index Terms—Agentic AI, Agents, Fintech, Artificial Intelligence

I. INTRODUCTION

Agentic AI-based autonomous systems have significant transformative potential for finance, particularly where traditional tech struggles with market volatility, risk assessment, and fraud detection. Agentic AI based systems offer compelling opportunities for automating intricate financial workflows such as portfolio management or loan processing, tasks that currently require critical human judgment despite their repetitive nature.

Beyond workflow automation, agentic AI techniques can streamline regulatory compliance and increase human expertise in complex areas such as risk modeling. Its impact extends to reshaping decision-making: consumers could gain personalized, data-driven financial guidance, while institutions benefit from faster risk assessments and portfolio adjustments.

Along with the aforementioned opportunities, there are several challenges toward widespread adoption of Agentic AI based techniques. In particular, we discuss ethical challenges that stem from fairness and transparency, challenges that arise from the evolving regulatory landscape and government frameworks, and secondary challenges that arise from privacy, security aspects, scalability, performance, and integration with legacy systems.

In this paper, we also discuss the various techniques and methods used to develop Agentic AI-based techniques. In particular, we also provide an overview of the design patterns (reflection, planning, tooling, collaboration) used for the development of such techniques. We considered the various components used in their inherent design (e.g., reasoning, memory, tooling) and provide an overview of how to train a large language model (LLM) to make use of tools and invoke the appropriate tool correctly.

We also discuss emerging research areas associated with Agentic AI. In particular, we discussed the opportunities and challenges associated with multi-model agents, challenges associated with human-agent collaboration, issues associated with learning and self-improvement, and the embodiment of Agents with the field of robotics and IoT. We also provided a brief overview of the risks associated with rapid development cycles (e.g., security vulnerabilities, prompt injections, data privacy, and government violations).

II. OPPORTUNITIES AND APPLICATION AREAS

In this section, we will emphasize the role of agentic AI in areas such as automation, decision making, personalized experience, security, and financial inclusion.

A. Automation

Agentic AI provides automation opportunities in the financial domain (for example, workflow management and compliance). Financial systems deal with complex workflows that currently require human intelligence to make choices that are hard to automate. For example, portfolio management, processing mortgage applications, credit card disputes, etc. currently require human intervention at several steps in the process. Many of these processes are repetitive with specific goals but require human intelligence in specific steps. In particular, Chan et al. [1] demonstrated how agentic AI techniques can play a useful role in the automation of such processes with examples of the Affirm, Upstart, and Monedo platforms.

In addition, Desai et al. [2] provided a survey of automation opportunities to automate the finance processing domain using generative AI. Further, Joshi et al. [3] proposed a robust agent framework, which combines front-end tools for user interaction, back-end models such as fine-tuned GPT architectures, and hybrid AI systems that integrate traditional econometric models with Gen AI.

Compliance is another important area in finance, where financial processes must be legally vetted. Agentic AI techniques can prove to be useful by proactively determining compliance requirements and reacting quickly to regulatory changes. Automation of legal processes is an interesting and evolving field of study and financial institutions will benefit from it [4].

B. Decision Making

Agentic AI techniques can also play a key role in the automation of institutional or consumer decision making in traditional and decentralized finance (DeFi).

1) *Decentralized Finance:* Agentic AI can make financial decisions or help the consumer in these decisions. Agents can work around the clock parsing a large amount of information focusing on consumer-provided goals [5]. Decentralized finance (DeFi) is an emerging technology in which traditional banking-like services are provided through public distributed ledgers similar to blockchains. The combination of agentic AI and DeFi is in its infancy, but we see it as one of the major areas of opportunity.

2) *Institutional decisions:* Financial institutions constantly make decisions about risk assessment, credit checks, portfolio adjustment to market conditions, etc. Any decision making that involves humans is time consuming and expensive. Agentic AI can help reduce the time that humans spend making these decisions.

Creditworthiness checks are an important part of any lending business. From simple credit cards to complex mortgage refinancing, some decisions are complex, involve a lot of communication with the consumer, and play a critical role in the profitability of the business. Rehman et al. [6] provide an overview of how AI can play a crucial role in rapid creditworthiness checks that allow new types of credits such as micro-financing. Most of the current research on AI and creditworthiness [7] focuses on just using AI, but agentic AI could be its natural extension.

C. Personalized experience

In recent years, financial services have rapidly migrated from static models to a more dynamic and personalized experience. This advancement can be further accelerated with Agentic AI given its capacity to exhibit problem-solving capabilities, autonomous initiative, and adaptability with limited supervision [8].

1) *Personalized financial advice and planning:* Agentic AI can act as personalized agents that allow customers to make well-informed decisions. These systems can analyze customer data, understand long- and short-term customer goals, and

provide customized recommendations on budgeting, savings, and investment strategies [2]. With Agentic AI's capability to adapt, these systems can learn user preferences and risk tolerance, making sure that the advice is aligned with the customer's risk profile [8].

2) *Portfolio management:* Agentic AI-powered tools can offer and execute customized investment strategies based on the customer's risk appetite, financial goal, and investment horizon [5]. Stock Agent is one such example that models investors' trading behavior in the stock market with dynamically adjusting strategies based on real-time stock market data.

D. Security

Financial systems deal with a vast amount of sensitive information and high-value transactions, making them a prime target for cyberattacks and financial fraud. With the financial sector being exponentially digitized in recent years, new vulnerabilities constantly emerge [9]. At the same time, customer expectations for data protection and financial regulations have also intensified. Agentic AI provides transformative approaches to these security challenges, as highlighted in the following.

1) *Improved Risk Assessment and Regulatory Compliance:* Agentic AI techniques can continuously monitor compliance by assessing whether financial activities align with complex and evolving regulatory requirements (e.g. GDPR, PSD2, AML laws), which with manual audits are time consuming [3]. Agentic AI systems by their nature of adaptability can analyze market movements and adjust risk scores in real time providing dynamic risk assessment. Agentic AI can also be helpful in a better credit card scoring and loan risk assessment by uncovering nuanced insights into customer behavior [2].

2) *Real-time Monitoring:* Agentic AI systems can proactively detect security threats by autonomously scanning and analyzing the vast amount of data without human supervision. Agentic AI systems can dynamically update the model for normal behavior and take actions when there is deviation. These systems can adjust strategies based on past and current data to improve trading outcomes in volatile markets [10].

E. Financial Inclusion

With Agentic AI's capabilities to assess nontraditional data such as customer behaviors, these systems can design personalized plans, microloans, or insurance products for under-represented customer segments. Such segments usually lack traditional credit histories, but with the help of Agentic AI techniques, custom yet reliable credit profiles can be built using mobile payment data, utility payments, and subscription data, thereby enabling services to be more inclusive. The scope of financial services can be further increased by Agentic AI's capabilities to autonomously communicate with customers in their preferred language and cultural context.

III. CHALLENGES

In order to utilize the true potential of Agentic AI and ensure its success, it is critical to address the following challenges.

A. Ethical Concerns

The integration of AI tools into financial analysis and decision-making requires a focus on protecting vulnerable populations from potential algorithmic and machine learning biases [11]. AI systems risk embedding bias and discrimination, resulting in unfair treatment for particular populations and eroding the foundations of equal and fair treatment. This is crucial, as financial systems often reflect and reinforce existing societal inequalities and economic disparities.

Ensuring fairness and transparency in AI-driven decision-making remains an important concern. For example, designing, developing, and deploying AI agents holistically in financial advice has ethical concerns. As researchers [12] pointed out, the design of AI financial advisors carries significant ethical risks. Issues such as potential discrimination, overly controlling (paternalistic) guidance, and unclear system objectives can lead to severe negative outcomes, potentially disadvantaging certain groups and damaging user confidence in algorithm-based financial recommendations.

B. Regulatory and Governance Challenges

The rise of agent-driven AI introduces several regulatory and governance challenges. There are no governance frameworks laid out by the regulators, which presents several challenges to organizations around the adoption of agentic AI.

1) *Evolving Regulatory Landscape*: Regulators are still grappling with how to govern AI, and the regulatory landscape is constantly evolving. Financial institutions need to find a way to stay ahead of these changes and ensure that their AI systems comply with all applicable regulations.

2) *Governance frameworks*: Establishing robust governance frameworks is essential to oversee the development, deployment, and use of agentic AI systems. This includes defining clear roles and responsibilities, setting ethical guidelines, and implementing monitoring and audit procedures.

C. Data Security and Privacy

Cyber threats, such as adversarial attacks and data breaches, pose significant risks to agentic AI systems.

1) *Data breaches*: Fan [13] talks about the implementation of LLM based techniques in the banking sector, which presents critical challenges regarding data privacy and security. Because these models process vast amounts of sensitive information, any data leakage or improper use exposes institutions to significant legal consequences and serious reputational harm, especially given the stringent global regulatory environment where data protection is paramount. The same risks apply towards the application of agentic AI systems. Encryption techniques are needed to ensure the confidentiality of the data. There could also be potential risks around data access management.

2) *Data privacy*: As the integration of agentic AI systems into the finance and banking sectors becomes more common, they pose additional challenges around data handling. Granting AI agents unfettered or nonanonymized access to databases increases the risk of exposing confidential information.

D. Technical Challenges

Agentic AI is still a developing technology, and there are multiple challenges to building robust scalable systems with good performance and seamless integrations with legacy systems.

1) *Lack of a standard development process*: Unlike traditional software, there is no standard development philosophy for agentic systems. Companies are still figuring out the best practices to integrate LLMs with available tools to solve tasks.

2) *Quality and availability of data*: Fine-tuning LLMs to use tools correctly requires high-quality training data. There are still challenges in ensuring data accuracy, completeness, and consistency, as well as accessing data from disparate sources.

3) *Integration with Legacy Systems*: Many financial institutions rely on outdated legacy systems that can be difficult to integrate with new AI technologies. Even without the challenges of legacy systems, there is a lack of a standard way for agents to connect to different data sources and APIs. Recently, the Model Context Protocol (MCP) [14] is gaining acceptance as a common standard, but the space remains largely fragmented.

4) *Scalability and performance*: Agentic AI systems need to be able to scale to handle large volumes of data and transactions, and they need to operate with high speed and reliability. The adoption of agentic AI systems is hindered by their huge computational and memory requirements, which poses challenges for deployment in resource-constrained environments. Chavan et al. [15] performed a cumulative analysis on the technical constraints posed by LLMs. For example, loading an LLaMa-70B model requires 140GB of VRAM excluding the memory required for model inferencing.

E. Operational Challenges

With the increasing adoption of agentic AI in industries like banking, finance & digital payments, there would be scalability issues around the adoption of complex agentic AI systems. In addition, to adopt agentic AI in the industry, there is a talent gap that would lead to difficulties in the robust implementation of agentic AI.

1) *Implementation Complexity*: With the financial markets already generating huge real-time data, generative AI models in order to improve accuracy of the recommendations carry out many simulations and hence further produce more synthetic datasets. Due to this, most agentic AI systems would have scalability issues [4].

2) *Transformation of the workforce*: As AI adoption increases, the nature of jobs in the financial sector would have to be transformed to adapt to the changes. There is a substantial skill gap that needs to be addressed in order for people to start utilizing agentic AI in financial industries in the right way. According to [16], a survey conducted in all organizations reported that 55% of the organizations surveyed felt that there was a skill gap. In addition, they also noted that there is a knowledge gap in organizations that is around

identifying the difference between agentic AI vs. generative AI [17].

IV. TECHNIQUES OR METHODOLOGIES

This section presents various methodologies used for the development of AI agents.

A. Design patterns for developing AI Agents

Andrew Ng in his lecture [18] proposed four key design patterns that underpin the development of AI agents and also represent an evolution from the simplest to the most sophisticated.

- 1) **Reflection:** This pattern focuses on self-improvement. Agents assess and refine their own output based on specific metrics or goals. It is like asking the model to "think twice" or "review its work."
- 2) **Tool Use :** In this pattern, agents are designed to use external tools or APIs to perform tasks that require specialized knowledge or capabilities, such as retrieving data, performing calculations, or interacting with other services.
- 3) **Planning:** This pattern is important for handling complex multistep tasks. Agents break down larger goals into smaller, manageable subtasks and create a strategy to achieve the final objective. Planning can often incorporate Tool Use and Reflection within individual steps.
- 4) **Multi-Agent Collaboration:** Multiple AI agents work together to solve a problem, where each agent may have specialized roles. This approach allows for the approach of more complex problems by making use of agents with specialized skills. This simulates a mixture of experts collaborating to solve a more general problem.

These four patterns are fundamental building blocks. In practice, an agentic AI system can combine multiple patterns. For example, a planning agent may use tools within its steps and make use of reflection to correct errors and to check the output at each intermediate step before proceeding. Similarly, multi-agent systems will have agents that individually use planning, tools, and reflection. Since the ReAct framework, first proposed by Yao et al., 2022 [19], interleaving reasoning and action has become a common pattern for agents. Yao et al. used the term "reason" to encompass both planning and reflection. At each step, the agent is asked to explain its thinking (planning), take actions, and then analyze observations (reflection), until the task is considered finished by the agent.

B. Key components of an agentic system

An agentic AI system can be broadly considered to be a combination of 3 key components: reasoning, memory, and external tools [20].

1) **Reasoning:** The reasoning module of an AI agent is a large language model. The LLM that handles reasoning, planning, and decision making. Modern implementations often use models like GPT-4o or specialized variants fine-tuned for agentic tasks. These models exhibit reasoning and planning capabilities through chain-of-thought reasoning that is useful for breaking down a high-level task into an executable plan.

2) **Memory:** Memory enables AI agents to retain, retrieve, and utilize information over time and overcomes a key challenge of LLMs - their stateless nature. The use of memory allows an agent to maintain context, store past interactions, maintain a coherent state, cache results of past queries, and remember user preferences. LLMs inherently do not have a memory and treat each interaction as a new one. Their memory is limited to the model's context length. Agentic frameworks equip foundational models with memory, allowing AI agents to store information beyond the context window of an LLM. Agentic memory can be classified into short-term, long-term, and episodic memory, each serving a different purpose [21].

- 1) **Short-term memory (STM)** enables an AI agent to recall recent interactions for immediate decision-making. This type of memory is useful in conversational AI, where maintaining context and a coherent state across multiple exchanges is required. STM is typically implemented using a rolling buffer or a context window, which holds a limited amount of recent data before being overwritten.
- 2) **Long-term memory (LTM)** is designed for permanent storage, implemented using databases, knowledge graphs, and vector embeddings. Also known as external knowledge in literature, long-term memory allows AI agents to store and recall information between different sessions. One of the most popular techniques for long-term memory is augmented retrieval generation (RAG) [22], where the agent retrieves relevant information from a stored knowledge base to respond to queries.
- 3) **Episodic Memory:** In addition to an external knowledge base for storing facts, agentic frameworks also store past interactions, known as episodes. Episodic memory is implemented by logging past queries, actions taken, key events, and their outcomes in a structured format [23]. Future tasks can retrieve episodes that are similar to the task at hand, based on relevance, and use the information to refine the planning process. Episodic memory allows AI agents to recall past interactions, which can be useful to learn from the past events and to make better decisions in the future.

3) **Tools :** Tools extend the capabilities of a large language model by allowing the agent to interact with the external world, access information beyond its training data cutoff, perform computations, and perform specialized tasks. The use of tools by AI agents is in the form of function calling with appropriate parameters and allows LLMs to interact with external services and APIs. For example, an LLM may invoke a weather API to get the weather of a location

```
weather_api.get_current_weather(  
    location="Munich, Germany",  
    unit="celsius"  
)
```

The tools used can be categorized into the following.

- **Information Retrieval:** Web search, databases, knowledge graphs, etc.

- Computation: Calculators or code interpreters.
- Execution: Updating a database entry, making a reservation, controlling devices, etc.

The next section explains how LLMs can be trained to use tools.

C. Training AI Agents to use tools

Along with memory, tool use is an active area of research for agentic AI, and several paradigms have been developed for both. This section provides a brief overview of how to train an LLM to make use of tools and invoke the appropriate tool correctly.

- Tool Definition: The agent needs to know which tools are available and how to use them. This is done by providing the LLM with detailed descriptions of the available tools. Such a description includes (1) Tool names, (2) Description of what the tool does, (3) Arguments that the tool expects, and (4) Format of the expected output/response.
- Tool Selection: The LLM must decide which tool is needed to perform the task. This decision is made either through chain-of-thought (CoT) reasoning, where the LLM reasons "I need to use finance tool to find the stock price of X" or through fine-tuning.
- Argument Generation: Once a tool is selected, the LLM must generate the correct arguments for that tool based on the current context and goal.
- Training LLMs to use the right tool: This can be achieved via zero-shot or few-shot prompting. Providing detailed instructions and few-shot examples of tool use within the context of LLM's prompt can enable LLM to use tools without explicit fine-tuning. Alternatively, fine-tuning can also be used [24] [25]. Models can be fine-tuned on datasets that contain examples of correct tool usage. These datasets contain examples like [input task, reasoning steps, API call syntax]. The objective of training is to train the LLM to generate syntactically correct API calls.

D. Evaluation of AI Agents

Evaluation of AI agents is still an open problem. Traditional machine learning metrics, for example, precision, recall, accuracy, are not sufficient for evaluating the performance of AI agents, and new metrics need to be devised. Similarly, only limited benchmarks exist for a small subset of tasks. At a high level, AI agents must be evaluated on their ability to complete a given within the given constraints. Some key evaluation metrics that have been proposed in the literature are as follows:

- 1) Task understanding: How well the agent understood the task and the intent of the user.
- 2) Task Success Rate: Measures the percentage of tasks completed correctly.
- 3) Reliability: Measures consistency across multiple trials of the same task.
- 4) Policy Compliance: Measures adherence to provided constraints. For example, a travel booking agent must find options only within the user's price range.

- 5) Resource Usage: How much memory, computation, and energy does the agent need to complete tasks?

E. AI Agent Frameworks

To enable the rapid development of AI agents, several frameworks have been introduced, each of which attempts to approach the problem in a slightly different way.

1) *LangChain*: LangChain [26] is designed to simplify the development of applications powered by large language models. Provides modular components for connecting LLMs to external tools such as databases and APIs, creating sequential workflows, and building agents that can perform tasks beyond simply generating text.

2) *LangGraph*: LangGraph [27] is an extension to LangChain, designed to build more complex AI workflows. LangGraph uses a graph-like structure to define how different AI agents or tools should work together to achieve an objective.

3) *AutoGen*: AutoGen [28], developed by Microsoft, is a framework for creating teams of AI agents that collaborate to solve problems. These agents can have different roles (eg a coder and a reviewer) and exchange messages with each other to complete tasks, such as debugging code or brainstorming ideas.

4) *CrewAI*: CrewAI [29] focuses on structured, role-based AI teams or "crews". It assigns specific roles, responsibilities, and tools to each agent. Agents can then work in sequence or in parallel to achieve a goal.

V. FORWARD OPINION ABOUT USAGE OF AGENTIC-AI

In this section, we discuss key emerging research areas that may be crucial for their future growth and some key risks associated with this rapid expansion.

A. Emerging Research Fields

- 1) Multimodal agent One of the emerging research trends is the development of a sophisticated multimodal agent. Future AI agents in the payment domain use multimodal data, memory mechanisms, reflection processes, and tool augmentation within a large language model framework to achieve significant performance improvements over existing methods [30]. The multimodal agent has the ability to perceive and interact with its environment using multiple data modalities and is not limited to text [31]. This capability allows them to handle more intricate and nuanced tasks compared to agents that process only textual information. The development of multimodal agents represents an evolution from language-only agents, enabling AI to understand and respond to a wider range of real-world information in a human-like way.
- 2) Human-Agent Collaboration As agents become more capable, they are very effectively acting as assistants through seamless human-agent interaction(HAI). Human-agent collaboration focuses on how humans and agents can work together effectively as a team to solve

a problem [32]. This is an evolution of multi-agent collaboration in which multiple AI agents work together to solve a problem [33]. Although significant research and development has already taken place in the field of multiagent collaboration, human-agent collaboration is still an emerging field of research.

- 3) **Learning, Adaptation, and Self-Improvement** One of the emerging research areas in the field of reinforcement learning of the AI agent is the development of an AI agent that can continuously learn from experience, adapt to changing environments, reflect on their performance, incorporate feedback effectively, and improve autonomously over time [34].
- 4) **Embodied Agentic AI** Integrating agentic AI capabilities with robotics and IoT devices is creating the rapidly emerging field of embodied agentic AI [35]. These systems can perceive, interact with, and act in the physical world [36]. While their direct role in core payments and finance is less common due to the digital nature of these domains, their potential in applications requiring physical interaction is vast. Futuristic use cases span autonomous scientific discovery and exploration, hyper-personalized healthcare and elder care companions, and complex operations in environments too hazardous or inaccessible for humans.

B. Risk With Rapid Expansion

- 1) **Rapid Development Cycles** With a high expansion rate, there is a risk of rapid development cycles of AI agents, which may not allow enough testing and validation. Increase the risk of unintended and potentially harmful behavior, especially for complex tasks or multi-agent systems where emergent behaviors are hard to predict. This could range from financial losses (e.g. incorrect transactions) to physical harm (in case of embodied AI agent) or critical system failures.
- 2) **Security Vulnerabilities** AI agents need vast amount of data and other APIs to function. This is especially true for multimodal AI agents that leverage multi-modal data sourced from different systems or environments. Rapid development can potentially cause overlooked security flaws, creating a new security vulnerability [37]. Deng et al. [38] has discussed some of the most common security vulnerabilities for AI agents:
 - **Prompt Injection** Tricking the agent into unintended actions [39].
 - **Data Ex-filtration** To perform complex tasks, the AI agent needs data access from multiple sources such as databases, APIs, etc. This broad data access makes them the potential medium for extracting data from various sources if compromised.
 - **Exploiting third party integrations** AI agent often has third-party integrations to perform complex tasks. Attackers can potentially exploit these third-party integrations if the AI agent is compromised.
- 3) **Ethical Oversights and Bias Amplification:** The rush to deploy may lead to inadequate vetting for biases embedded in training data or algorithms. Agentic AI could then autonomously perpetuate or even scale discriminatory practices [40] in areas like hiring, loan applications, or content management before these issues are fully understood or corrected.
- 4) **Data privacy violations** Agents often require access to large amounts of data to function effectively. Rapid integration across systems increases the risk that sensitive personal or proprietary data could be misused or processed in non-compliant ways (e.g., violating GDPR, CCPA) [41].
- 5) **Regulatory Lag and Governance Gaps** Technology development, especially at a rapid pace, often exceeds the ability of regulators and policy makers to create effective governance frameworks. This can lead to a period where powerful agentic AI systems operate in a grey area without clear rules, standards, or accountability structures [42].

VI. CONCLUSION

In this paper, our objective is to explore the emerging field of agentic AI in finance by exploring various opportunities through reshaping institutional and personalized decision making, automating financial workflows, and streamlining regulatory compliances. We also present a brief overview of challenges associated via government frameworks, integration with existing systems, and privacy / security standards.

We also discuss various methodologies along with the underlying design patterns and composite structures required to build such systems. We also discussed available frameworks for building and escape such systems. More importantly, we discuss the future of this area within the realm of multi-model agents, human-agent collaboration, and issues associated with integration with secondary sources.

To the best of our knowledge, this is the first work to comprehensively summarize the usage of Agentic-AI techniques in the finance domain. This summarization would provide a clear overview where future efforts can be prioritized, potentially accelerating the development and adoption of agent-based AI techniques in finance. In addition, this analysis could also facilitate the introduction of innovative ideas, thereby benefiting a wider range of industries and applications.

ACKNOWLEDGMENT

We would like to express gratitude to Ashish Gupta & Prateek Dudeja for reviewing the paper.

REFERENCES

- [1] L. Chan, L. Hogaboam, and R. Cao, "Artificial intelligence in credit, lending, and mortgage," in *Applied Artificial Intelligence in Business: Concepts and Cases*. Springer, 2022, pp. 201–211.
- [2] A. P. Desai, T. Ravi, M. Luqman, G. Mallya, N. Kota, and P. Yadav, "Opportunities and challenges of generative-ai in finance," in *2024 IEEE International Conference on Big Data (BigData)*, 2024, pp. 4913–4920.

- [3] S. Joshi, "Gen ai for market risk and credit risk learn agentically powered gen ai; gen ai agent framework for financial risk management," *Gen AI Agent Framework for Financial Risk Management (January 15, 2025)*, 2025.
- [4] C. G. O'Grady and C. OG, "Agentic workflows in the practice of law—ai agents as ethics counsel," *Arizona Legal Studies Discussion Paper*, pp. 25–03, 2024.
- [5] D. B. Acharya, K. Kuppan, and B. Divya, "Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey," *IEEE Access*, 2025.
- [6] M. A. Rehman, M. Ahmed, and S. Sethi, "Ai-based credit scoring models in microfinance: Improving loan accessibility, risk assessment, and financial inclusion," *The Critical Review of Social Sciences Studies*, vol. 3, no. 1, pp. 2997–3033, 2025.
- [7] E. Marevac, S. Patković, and E. Zunić, "Decision-making ai for customer worthiness and viability," in *2023 22nd International Symposium INFOTEH-JAHORINA (INFOTEH)*. IEEE, 2023, pp. 1–6.
- [8] H. Clatterbuck, C. Castro, and A. M. Morán, "Risk alignment in agentic ai systems," 2024.
- [9] A. P. Desai, T. Ravi, M. Luqman, M. Sharma, N. Kota, and P. Yadav, "Gen-ai for user safety: A survey," in *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 2024, pp. 5315–5324.
- [10] D. B. Acharya, K. Kuppan, and B. Divya, "Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey," *IEEE Access*, vol. 13, pp. 18912–18936, 2025.
- [11] A. Wagstaff, "Reflections on and alternatives to who's fairness of financial contribution index," *Health economics*, vol. 11, no. 2, pp. 103–115, 2002.
- [12] L. Bruggen, R. Gianni, F. de Haan, J. Hogreve, D. Meacham, T. Post, and M. van der Werf, "Ai-based financial advice: An ethical discourse on ai-based financial advice and ethical reflection framework," *Journal of Public Policy & Marketing*, vol. 0, no. 0, p. 07439156241302279, 0.
- [13] M. Fan, "Llms in banking: Applications, challenges, and approaches," in *Proceedings of the International Conference on Digital Economy, Blockchain and Artificial Intelligence*, ser. DEBAI '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 314–321.
- [14] Anthropic, "Model context protocol," accessed April 30, 2025.
- [15] A. Chavan, R. Magazine, S. Kushwaha, M. Debbah, and D. Gupta, "Faster and lighter llms: A survey on current challenges and way forward," *arXiv preprint arXiv:2402.01799*, 2024.
- [16] P. Sawant, "Agentic ai: A quantitative analysis of performance and applications," *Preprints*, February 2025.
- [17] L. Ackerman, "Perceptions of agentic ai in organizations: Implications for responsible ai and roi," 2025.
- [18] <https://www.facebook.com/DeepLearningAIHQ/>, "One Agent For Many Worlds, Cross-Species Cell Embeddings, and more — deeplearning.ai," [Accessed 25-04-2025].
- [19] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafraan, K. Narasimhan, and Y. Cao, "React: Synergizing reasoning and acting in language models," in *International Conference on Learning Representations (ICLR)*, 2023.
- [20] P. Zhao, Z. Jin, and N. Cheng, "An in-depth survey of large language model-based artificial intelligence agents," *arXiv preprint arXiv:2309.14365*, 2023.
- [21] Z. Zhang, X. Bo, C. Ma, R. Li, X. Chen, Q. Dai, J. Zhu, Z. Dong, and J.-R. Wen, "A survey on the memory mechanism of large language model based agents," *arXiv preprint arXiv:2404.13501*, 2024.
- [22] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel *et al.*, "Retrieval-augmented generation for knowledge-intensive nlp tasks," *Advances in neural information processing systems*, vol. 33, pp. 9459–9474, 2020.
- [23] W. Zhong, L. Guo, Q. Gao, H. Ye, and Y. Wang, "Memorybank: Enhancing large language models with long-term memory," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 17, 2024, pp. 19724–19731.
- [24] Q. Xu, F. Hong, B. Li, C. Hu, Z. Chen, and J. Zhang, "On the tool manipulation capability of open-source large language models," *arXiv preprint arXiv:2305.16504*, 2023.
- [25] T. Schick, J. Dwivedi-Yu, R. Dessì, R. Raileanu, M. Lomeli, E. Hambro, L. Zettlemoyer, N. Cancedda, and T. Scialom, "Toolformer: Language models can teach themselves to use tools," *Advances in Neural Information Processing Systems*, vol. 36, pp. 68539–68551, 2023.
- [26] H. Chase, "LangChain," Oct. 2022.
- [27] L. AI, "Langgraph," 2022.
- [28] Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, L. Jiang, X. Zhang, S. Zhang, J. Liu *et al.*, "Autogen: Enabling next-gen llm applications via multi-agent conversation," *arXiv preprint arXiv:2308.08155*, 2023.
- [29] C. Inc., "crewai," 2024.
- [30] Z. Durante, Q. Huang, N. Wake, R. Gong, J. S. Park, B. Sarkar, R. Taori, Y. Noda, D. Terzopoulos, Y. Choi, K. Ikeuchi, H. Vo, L. Fei-Fei, and J. Gao, "Agent ai: Surveying the horizons of multimodal interaction," 2024.
- [31] J. Xie, Z. Chen, R. Zhang, X. Wan, and G. Li, "Large multimodal agents: A survey," *arXiv preprint arXiv:2402.15116*, 2024.
- [32] R. K. Bellamy, S. Andrist, T. Bickmore, E. F. Churchill, and T. Erickson, "Human-agent collaboration: Can an agent be a partner?" in *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, 2017, pp. 1289–1294.
- [33] X. Feng, Z.-Y. Chen, Y. Qin, Y. Lin, X. Chen, Z. Liu, and J.-R. Wen, "Large language model-based human-agent collaboration for complex task solving," 2024.
- [34] P. Ponnusamy, A. Ghias, Y. Yi, B. Yao, C. Guo, and R. Sarikaya, "Feedback-based self-learning in large-scale conversational ai agents," *AI magazine*, vol. 42, no. 4, pp. 43–56, 2022.
- [35] J. Duan, S. Yu, H. L. Tan, H. Zhu, and C. Tan, "A survey of embodied ai: From simulators to research tasks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 6, no. 2, pp. 230–244, 2022.
- [36] J. N. Bailenson, K. Swin, C. Hoyt, S. Persky, A. Dimov, and J. Blascovich, "The independent and interactive effects of embodied-agent appearance and behavior on self-report, cognitive, and behavioral markers of copresence in immersive virtual environments," *Presence*, vol. 14, no. 4, pp. 379–393, 2005.
- [37] Y. He, E. Wang, Y. Rong, Z. Cheng, and H. Chen, "Security of ai agents," 2024.
- [38] Z. Deng, Y. Guo, C. Han, W. Ma, J. Xiong, S. Wen, and Y. Xiang, "Ai agents under threat: A survey of key security challenges and future pathways," *ACM Computing Surveys*, vol. 57, no. 7, pp. 1–36, 2025.
- [39] Y. Liu, G. Deng, Y. Li, K. Wang, Z. Wang, X. Wang, T. Zhang, Y. Liu, H. Wang, Y. Zheng *et al.*, "Prompt injection attack against llm-integrated applications," *arXiv preprint arXiv:2306.05499*, 2023.
- [40] G. B. Mensah, "Artificial intelligence and ethics: a comprehensive review of bias mitigation, transparency, and accountability in ai systems," *Preprint, November*, vol. 10, no. 1, 2023.
- [41] J. M. Such, A. Espinosa, and A. García-Fornes, "A survey of privacy in multi-agent systems," *The Knowledge Engineering Review*, vol. 29, no. 3, pp. 314–344, 2014.
- [42] N. A. Smuha, "From a 'race to ai' to a 'race to ai regulation': regulatory competition for artificial intelligence," *Law, Innovation and Technology*, vol. 13, no. 1, pp. 57–84, 2021.