

A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist

Wentao Zhang
Nanyang Technological University
Singapore
wt.zhang@ntu.edu.sg

Jiaze Sun
National University of Singapore
Singapore
e0564914@u.nus.edu

Xinyu Cai
Longtao Zheng
Nanyang Technological University
Singapore
{xinyu009,longtao001}@e.ntu.edu.sg

Lingxuan Zhao*
Haochong Xia*
Nanyang Technological University
Singapore
{zhao0375,haochong001}@e.ntu.edu.sg

Molei Qin
Xinyi Li
Yuqing Zhao
Nanyang Technological University
Singapore
{molei001,lix0067,ZHAO0348}@e.ntu.edu.sg

Xinrun Wang[†]
Nanyang Technological University
Singapore
xinrun.wang@ntu.edu.sg

Shuo Sun
Nanyang Technological University
Singapore
shuo003@e.ntu.edu.sg

Yilei Zhao
Zhejiang University
China
yilei_zhao@zju.edu.cn

Bo An[†]
Nanyang Technological University
Skywork AI
Singapore
boan@ntu.edu.sg

ABSTRACT

Financial trading is a crucial component of the markets, informed by a multimodal information landscape encompassing news, prices, and Kline charts, and encompasses diverse tasks such as quantitative trading and high-frequency trading with various assets. While advanced AI techniques like deep learning and reinforcement learning are extensively utilized in finance, their application in financial trading tasks often faces challenges due to inadequate handling of multimodal data and limited generalizability across various tasks. To address these challenges, we present FinAgent¹, a multimodal foundational agent with tool augmentation for financial trading. FinAgent’s market intelligence module processes a diverse range of data—numerical, textual, and visual—to accurately analyze the financial market. Its unique dual-level reflection module not only enables rapid adaptation to market dynamics but also incorporates a diversified memory retrieval system, enhancing the agent’s ability to learn from historical data and improve decision-making processes. The agent’s emphasis on reasoning for actions fosters

trust in its financial decisions. Moreover, FinAgent integrates established trading strategies and expert insights, ensuring that its trading approaches are both data-driven and rooted in sound financial principles. With comprehensive experiments on 6 financial datasets, including stocks and Crypto, FinAgent significantly outperforms 12 state-of-the-art baselines in terms of 6 financial metrics with over 36% average improvement on profit. Specifically, a 92.27% return (a 84.39% relative improvement) is achieved on one dataset. Notably, FinAgent is the first advanced multimodal foundation agent designed for financial trading tasks.

CCS CONCEPTS

• Information systems → Data mining; • Computing methodologies → Machine learning; • Applied computing → Electronic commerce.

KEYWORDS

Large Language Models, Quantitative Trading, Financial AI Agents

ACM Reference Format:

Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, Longtao Zheng, Xinrun Wang, and Bo An. 2024. A Multimodal Foundation Agent for Financial Trading: Tool-Augmented, Diversified, and Generalist. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3671801>

1 INTRODUCTION

Financial markets are essential for economic stability, facilitating capital allocation and risk management. Financial trading systems, developed from technical analysis strategies [9], enhance these

*Lingxuan Zhao and Haochong Xia contributed equally to this research.

[†]Corresponding Authors.

¹The full technical report is available at <https://arxiv.org/abs/2402.18485>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0490-1/24/08

<https://doi.org/10.1145/3637528.3671801>




Figure 1: Overview of FinAgent.

markets by enabling efficient trading. Rule-based trading systems are rigid and struggle to adapt to market volatility, often resulting in underperformance in evolving markets. Reinforcement learning-based systems[2] demonstrate enhanced adaptability but encounter substantial obstacles, such as the need for extensive training data and the inexplainsibility of decision-making processes. Additionally, they struggle with generalizing across diverse market conditions, are sensitive to market noise, and often fail to integrate multimodal market intelligence like news and reports into their analysis. The financial trading landscape demands more advanced machine-learning methods to address complex market dynamics, seeking to move beyond the limitations of rule-based and RL methods.

Recently, Large Language Models (LLMs) have showcased their potential in a range of decision-making tasks when applied in AI agents [28, 35, 43, 57], marking a significant expansion beyond natural language processing into more complex, task-specific functions. This advancement includes the integration of memory and planning modules, which enable these agents to adapt within dynamic environments, akin to human cognitive processes. This evolution has been further pushed by the advent of multimodal LLMs like GPT-4V [26], which enhances the capabilities of LLMs by processing both textual and visual data. Moreover, the integration of tool-augmented models like Toolformer [33] empowers LLMs to utilize external tools, thus elevating their decision-making abilities in complex scenarios. This combination of adaptability and enhanced processing capabilities offers new possibilities in fields such as fintech, where nuanced analysis and adaptation are important.

LLMs have demonstrated remarkable capabilities in analyzing and interpreting financial data, as evidenced by developments like BloombergGPT [48], and FinGPT [50]. However, there is a natural gap between QA tasks and sequential decision-making in trading. Although FinMEM [56] is an LLM trading agent with a human-aligned memory mechanism and character design, the full capabilities of LLMs as comprehensive autonomous trading systems remain underexplored, particularly in their ability to interpret multimodal data and utilize diverse tools. The challenges in navigating the complexities of financial markets are identified as follows:

- **Ch1: Insufficient Multimodal Data Processing Ability.** Processing numerical, textual, and visual market intelligence data significantly requires advanced analytical methods to extract key insights and predict market trends.
- **Ch2: Imprecise information retrieval.** Mixing retrieval with main tasks and relying on brief summaries causes imprecise searches, introducing irrelevant data and reducing performance.
- **Ch3: Adaptability in Rapidly Evolving Markets.** Financial trading requires the ability to quickly adapt to fluctuating market conditions. Traditional methods often fall short, highlighting the necessity for models capable of responding to real-time data and adjusting strategies according to historical market trends.

- **Ch4: Integration of Domain Knowledge.** Current models often struggle to integrate established methods such as expert guidance and advanced trading tools effectively, leading to a decline in both the effectiveness and depth of market analysis.
- **Ch5. Reasoning for Actions.** The black-box nature of many sophisticated AI models, directly giving results of decisions without providing the reasoning process.

To address the challenges of adapting the multimodal LLMs to the dynamic and information-rich financial trading tasks, we present FinAgent, a multimodal foundation agent that integrates both textual and visual information for a comprehensive analysis of market dynamics and historical trading patterns. Specifically, FinAgent’s market intelligence module processes multimodal data, such as numerical, textual, and visual, to provide precise analysis of financial market trends, offering insights for future trading tasks (**Ch1**). A uniquely designed dual-level reflection module is developed, capable of not only rapidly adapting to market dynamics but also enhancing the agent’s ability to learn from historical data and improve its decision-making process (**Ch2**). FinAgent introduces a diversified memory retrieval system for the market intelligence and reflection modules, separating trading and retrieval tasks to enhance focus on their specific functions and minimize noise in the results (**Ch3**). Finally, the decision-making module incorporates expert knowledge, comprising both supplementary expert guidance and auxiliary expert strategies, to guide the agent’s decisions. This emphasis on providing reasoned explanations for actions fosters trust in its financial decisions (**Ch4 & Ch5**). Specifically, our contributions are four-fold:

- We introduce the market intelligence module, which is able to extract key insights from multimodal datasets encompassing asset prices, visual representations, news, and expert analyses, offering a multifaceted view across various markets.
- We not only generate summaries for trading tasks but also provide query fields for retrieval tasks. These query texts include different retrieval types, tailored to enable focused retrieval of specific types of information.
- Our dual-level reflection module combines a low-level reflection that analyzes market price movement for insights, while the high-level reflection assesses past trading decisions for improvement, emulating the learning process in decision-making.
- We employ a suite of tools in FinAgent, including expert guidance and technical indicator-based advanced trading strategies, to infuse domain knowledge in financial trading.

With comprehensive experiments on 6 financial datasets, including stocks and Crypto, FinAgent significantly outperforms 12 state-of-the-art baselines in terms of 6 financial metrics with over 36% average improvement on profit. Specifically, a 92.27% return (a 84.39% relative improvement) is achieved on one dataset. Notably, FinAgent is the first advanced multimodal foundation agent designed for financial trading tasks.

2 RELATED WORK

2.1 LLM Agents for Decision Making

The field of artificial intelligence and natural language processing has reached a significant milestone with the emergence of LLMs

Table 1: Comparison of FinAgent versus trading strategies and LLM agents. Brief introduction can be found in Section 5.3.

Method	Market Intelligence				Tool Use		Inference & Extension				
	News	Reports	Price	Visual Data	Info	Tools	Preference	Training Scheme	Planning	Explainability	Generalization
Rule-based	✗	✗	✓	✗	✗	✗	✗	Hyper-parameter Tuning	Myopic	-	Single trading task
RL method	✗	✗	✓	✗	✗	✗	✗	Model training	Sequential	✗	Single trading task
FinGPT	✓	✗	✓	✗	✗	✗	✗	LLM Fine-tuning	Myopic	✓	Limited trading tasks
FinMem	✓	✓	✓	✗	✗	✗	✓	Reflection	Myopic	✓	Multiple trading tasks
FinAgent	✓	✓	✓	✓	✓	✓	✓	Reflection	Sequential	✓	Multiple trading tasks

like ChatGPT [24] and GPT-4 [25]. BloombergGPT [48] introduced the first LLM in the finance domain, combining financial and text data, but without public access. FinGPT [50] proposed the first open-source finance LLMs, incorporating reinforcement learning with human feedback.

While LLMs achieve impressive performance in NLP tasks [5, 42], more works explored the capability of LLMs to function not just as language processors but as agents capable of performing complex tasks. Initiatives like AutoGPT [51] and MetaGPT [12], Voyager [43], and AI agents [28, 35] expand LLMs’ capabilities to complex tasks involving reasoning and collaboration, significantly advancing technology and impacting daily life. FinMEM [56] presents an LLM agent with a human-aligned memory mechanism and character design for automated trading.

Recently, there has been growing interest in enhancing LLM agents with external tools and modular methods as AI agents. Tool-augmented Language Models (TALM) [22, 27, 33, 41] have been evaluated through recent benchmarks, such as ScienceQA and TabMWP [4, 17, 18, 20, 36, 46], designed to assess their ability to tackle intricate reasoning challenges, particularly those requiring the use of external tools. These improvements enable LLMs to retrieve current information through web searches [22] and to apply specialized knowledge from external sources [55].

However, a major limitation of LLM agents is their dependence on text-based information, which limits their perception and interaction with the environment. Introducing models equipped with vision capabilities, such as the latest iteration of GPT-4V [26], marks a pivotal breakthrough. There has also been the emergence of multimodal agents [19, 53, 57] utilizing the visual capabilities of multimodal large language models to perform tasks previously unachievable by text-only agents. Most existing LLMs in finance focus on NLP tasks, and their potential in trading is not fully explored. FinAgent is a multi-modal, tool-augmented LLM foundation agent for financial trading to bridge the gap.

2.2 AI for Financial Trading

AI techniques have been widely used in various financial trading tasks. RNN-based such as GRU [23] and LSTM [44] models are popular for stock prediction since they are specifically designed to capture temporal patterns in sequential data. Another direction of work employs graph-based DL models to model pair-wise relations between stocks. For instance, Feng et al. [10] enhance graph convolutional networks (GCNs) with temporal convolutions for mining inter-stock relations. Sawhney et al. [30] focus on stock industry data and links between company CEOs. Tree-based models [14] also achieve robust performance. Xu and Cohen [49] propose a

variational autoencoder architecture to extract latent information from tweets. Chen et al. [3] enhance trading strategy design with the investment behaviors of professional fund managers. Other data sources such as economics news [13] and earning calls [31] are also used to improve the prediction performance. Sun et al. [40] introduce a novel three-stage ensemble learning method. Reinforcement learning [39] has achieved success in finance with algorithms, platform [38], and evaluation toolkits [37]. However, most of these methods are hindered by their focus on price data and limited generalization, necessitating advanced techniques that can integrate multimodal intelligence and navigate complex market dynamics.

3 PROBLEM FORMULATION

We first introduce the Markov Decision Process (MDP) formulation of financial trading. Later on, we provide the formal formulation of FinAgent, which integrates LLMs into the RL pipeline to enable flexible reasoning and decision-making in financial trading.

3.1 Financial Trading as MDP

A financial trading task involves sequentially making investment decisions (e.g., buy/sell stocks) to maximize total profit under certain risk tolerance [39]. We formulate it as an MDP under a classic RL scenario following [15, 38], where an agent (investor) interacts with an environment (the financial market) to make actions (investment decisions) at discrete time to earn rewards (profits). The MDP is constructed by a 5-tuple $(S, \mathcal{A}, \mathcal{T}, R, \gamma)$. Specifically, S is a finite set of states. \mathcal{A} is a finite set of actions. The state transition function $\mathcal{T} : S \times \mathcal{A} \times S \rightarrow [0, 1]$ encapsulates transition probabilities between states based on chosen actions. The reward function $R : S \times \mathcal{A} \rightarrow R$ quantifies the immediate reward of taking an action in a state. The discount factor is $\gamma \in [0, 1)$. A policy $\pi : S \times \mathcal{A} \rightarrow [0, 1]$ assigns each state $s \in S$ a distribution over actions, where $a \in \mathcal{A}$ has probability $\pi(a|s)$. During training, the agent is in charge of making investment decisions at each time step through one whole trading period and tries to learn an optimal policy (investment strategy) that maximizes the expected sum of discounted reward (overall profit): $\pi_{\theta^*} = \arg \max_{\pi_{\theta}} \mathbb{E}_{\pi_{\theta}} [\sum_{i=0}^T \gamma^i r_{t+i} | s_t = s]$.

Specifically, we focus on single asset (e.g., stock or Crypto) trading. A *state* represents RL agents’ perception on the financial market based on price information, limited order book [29], technical indicators, trend prediction [54], financial news [32], experts’ investment behaviors [8] and overall market status [47]. The *action space* includes three choices to buy, sell or hold the asset [7, 16]. The *reward function* leverages the change of market capitals (earned/lost money) [16] with consideration of commission fee [38, 45].

3.2 Problem Formulation

We further integrate multimodal LLMs into the RL framework [6], enabling the flexible definition of the reasoning processes. In FinAgent formulation, we focus on the necessity of defining, learning, and applying these processes independently. We extend the classic RL optimization problem for FinAgent as follows:

$$\pi_{\theta^*} = \arg \max_{\pi_{\theta}} \mathbb{E}_{\pi_{\theta}} \left[\sum_{i=0}^T \gamma^i r_{t+i} | s_t = s, \mu_t = \mu \right], \quad (1)$$

where r_t is the reward at the time step t that depends on the environmental state s_t and action a_t . $\mu(\cdot)$ are specialized modules that encapsulate beneficial internal reasoning processes. Note that a state contains multimodal information including textual, numerical, and visual data. Faced with a task λ and equipped with a memory Mem_t^λ and a tool $Tool_t^\lambda$, FinAgent acting as the multimodal LLM agent, determines its action a_t through the following process:

$$\begin{aligned} \pi_{\text{FinAgent}}(a_t | s_t, \mu_t) &\equiv \mathcal{D}^\lambda \left(LLM \left(\phi_D^\lambda(s_t, \mu_t) \right) \right) \\ \mu_t &= \mu(s_t, Mem_t^\lambda, Tool_t^\lambda) \end{aligned} \quad (2)$$

where $\phi(\cdot)$ is a task-relevant prompt generator. The prompt is then passed to a multimodal LLM, from which a response is generated. Finally, the response is parsed through the task-specific action parsing function $\mathcal{D}^\lambda(\cdot)$ to perform compatible actions in the environment.

FinAgent is a multimodal LLMs agent in this framework specifically designed for financial trading, which contains five core modules, namely market intelligence module (M), memory module (Mem), low-level reflection module (L), high-level reflection module (H) and decision-making module (D). We can define the μ_t and other modules as follows:

$$\begin{aligned} \mu_t &= \mu(s_t, Mem_t^\lambda, Tool_t^\lambda) = \mu(M_t^\lambda, L_t^\lambda, H_t^\lambda, Tool_t^\lambda) \\ M_t^\lambda &= LLM(\phi_M^\lambda(s_t, Mem_t^{M,\lambda})) \\ L_t^\lambda &= LLM(\phi_L^\lambda(M_t^\lambda, KC_t, Mem_t^{L,\lambda})) \\ H_t^\lambda &= LLM(\phi_H^\lambda(M_t^\lambda, TC_t, Mem_t^{H,\lambda})), \end{aligned} \quad (3)$$

where M , Mem , L , H , D correspond to each module respectively, $Mem^{*,\lambda}$ denotes the memory of M , L , and H . KC and TC represent the Kline chart and Trading chart. ϕ_*^λ denotes the prompt generator corresponding to each module associated with task λ .

Therefore, with the integration of memory mechanism, augmented tools, and several designed modules, the overall objective of FinAgent is to find policies as described in Eq. (2) to optimize total discounted returns:

$$\begin{aligned} \pi_{\text{FinAgent}}^* &= \arg \max_{\pi(\cdot), \mu(\cdot)} \mathbb{E}_{\pi} \left[\sum_{i=0}^T \gamma^i r_{t+i} | s_t = s, \mu_t = \mu \right] \\ \text{s.t. } \pi(a_t | s_t, \mu_t) &= \mathcal{D}^\lambda \left(LLM \left(\phi_D^\lambda(s_t, \mu_t) \right) \right) \text{ with Eq.(3)} \quad \forall t. \end{aligned} \quad (4)$$

4 FINAGENT FRAMEWORK

As shown in Figure 3, the FinAgent framework comprises five core modules. Specifically, the market intelligence module (§4.1) is responsible for collecting, collating, summarizing, and analyzing market information, which includes daily updates on stock news, prices, and monthly and quarterly financial reports. The low-level reflection module (§4.3) establishes the inherent correlation between market intelligence and price changes. And the high-level reflection module (§4.3) involves reflecting on market conditions, price changes, and other factors in the context of outcomes from past trading decisions, which aims to derive insights from previous experiences and identify potential improvement in profitability by

assessing the efficacy of historical decisions and offering recommendations for future decision-making processes. The primary role of the memory module (§4.2) is to support the aforementioned three modules by offering storage capabilities and vector retrieve functions. The tool-augmented decision-making module (§4.4) integrates the aforementioned information, along with augmented tools and trader preferences, to make final investment decisions with a comprehensive analysis.

4.1 Market Intelligence Module

To make profitable investment decisions, it is beneficial to collect, summarize, analyze, and extract key insights from various multimodal financial data sources. We design the market intelligence module to achieve this goal. Market intelligence typically involves daily data about the macro environment, current market conditions or investors' sentiments that inform investment and trading decisions. In FinAgent, we harness the power of both the latest and historical news, financial reports, and asset prices related to the targeted asset in order to inform and optimize trading decisions.

Latest Market Intelligence. This module mainly consists of asset news and daily asset prices. However, it is not confined to these elements alone. Any information impacting the market can be encompassed within our framework as part of the latest market intelligence. The objective of this component is to evaluate the sentiment² of each market intelligence item regarding its influence on future asset prices and to provide a detailed summary of whether the market has recently exhibited bearish or bullish tendencies, thereby assisting in informed decision-making.

Nevertheless, historical data can offer insights into patterns that might influence future pricing and potentially affect current and upcoming market dynamics. For instance, if a past product launch significantly boosted a company's stock, a recent launch might have a similar effect³. We hope to incorporate these historical experiences and patterns into FinAgent's considerations. This inspired us to add two additional functional layers: retrieving relevant information from past market intelligence and summarizing key insights and historical experiences from them.

Diversified Retrieval Operation. A straightforward approach involves using the summary of the latest market intelligence as the query text and then employing an LLM to extract its semantically rich embeddings. This allows for retrieving past market intelligence with similar content through vector similarity. However, adopting this approach inevitably comes with two significant shortcomings: i) the summary of recent market intelligence is primarily aimed at supporting subsequent trading decision tasks, not for retrieval tasks. The significant gap between these two objectives can lead to unsatisfactory retrieval results; ii) some noise unrelated to the retrieval task may be contained in the summary, directly affecting the retrieval results. To address these challenges, diversified retrieval is implemented in FinAgent. Specifically, we have introduced an additional query text field to the output of the latest market intelligence component, which is dedicated to serving retrieval tasks in parallel with the summary that caters to trading tasks. It is worth

²Market intelligence can be categorized as positive, negative, or neutral based on its impact on market perceptions and potential outcomes.

³Some news will detail the percentage increase or decrease in a company's stock price after some event occurs.




Figure 2: The overall architecture of FinAgent. The ordinal numbers in the figure represent the order of execution, where augmented tools are implemented with the decision-making module.

emphasizing that we can define various retrieval types⁴ to enable an agent to retrieve past market intelligence from multiple perspectives, in multiple senses, and with a purpose. As shown in Figure 3, there are M retrieval types, so retrieving top K historical market intelligence separately can form a combination of $M \times K$ market intelligence in the past. This approach assigns specific retrieval types to each piece of historical information accompanying the summaries. This nuanced labeling facilitates a more targeted and efficient search and retrieval process.

Past Market Intelligence. Once similar past market intelligence is searched, it undergoes the summarising step, delivering key insights tailored to augment trading decisions. This meticulous approach ensures that only the most relevant information is incorporated, mitigating the impact of noise and maximizing the utility of historical data in informing trading strategies.

4.2 Memory Module

The memory mechanism [6, 28, 57] is crucial in LLM Agents for effectively handling extensive texts, grasping the context, ensuring the coherence of conversations, and improving the agent’s comprehension and logical abilities. In the context of multimodal LLM agents for financial trading, memory mechanisms play a crucial role in three main aspects: i) *Acuity*. This feature enables multimodal LLM agents to use market news, financial reports, and other information for better market forecasting. By analyzing historical data and current events, these agents can predict market trends and asset prices more accurately, aiding in effective trading decisions. ii) *Adaptability*. As market conditions change rapidly, memory mechanisms allow multimodal LLM agents to quickly learn and adapt. By continuously analyzing market data and trading outcomes, these agents adjust their strategies to handle volatility and seize new opportunities. iii) *Amendability*. It helps multimodal LLM agents learn from past mistakes and successful trades. By reflecting on these experiences, agents can avoid repeating errors and improve

their trading strategies. This continuous learning enhances their performance and creates more robust, efficient trading strategies.

To realize the 3A superiority - *Acuity*, *Adaptability*, and *Amendability* - in the memory mechanism, our development of the memory module employed a vector storage architecture. This module is composed of three main components: market intelligence memory (service for (§4.1)), low-level reflection memory (service for (§4.3)), and high-level reflection memory (service for (§4.3)). As shown in Figure 3, the summarize operation creates a query text field for each module, enhancing memory storage and retrieval. The market intelligence module uniquely retrieves past data through query text, using vector representations for efficient matching based on the vector similarity. All analyses and summaries from the market intelligence, low-level reflection, and high-level reflection modules are stored in the memory module. This integration equips the agent with extensive market data and insights, improving its decision-making capabilities.

4.3 Reflection Module

A reflection module is incorporated into the agent’s design to emulate the cognitive learning process inherent in human decision-making. The reflection framework is divided into low-level reflection and high-level reflection, each serving distinct purposes to enhance the agent’s trading decisions. The low-level reflection module involves reflecting on the relationship between the agent’s observations (e.g., news, financial reports, Kline chart and technical indicators) and the resultant price movements in the market, drawing connections between the provided information and the actual price changes. Whereas the high-level reflection step examines past decisions, tracking both the agent’s actions and the subsequent price movements in order to learn from past successes or mistakes. **Low-level Reflection Module** The primary focus of the low-level reflection module is to analyze the connection between the given market intelligence together with the Kline chart and technical indicators and past and future price changes to enhance decision-making. After taking in the price change data, the module generates detailed analysis for varying temporal horizons, spanning short-term, medium-term to long-term perspectives. The emphasis is

⁴The retrieval types include short-term, medium/long-term market impacts, asset price increase/decrease, market trends bearish/bullish, news/reports, etc.




Figure 3: Case studies of FinAgent. We only display the partial prompt for brevity.

placed on identifying potential patterns in the price movements of the targeted stock and deriving insights from how the given market intelligence summaries and Kline chart analysis can lead to such price movements. In order to facilitate future access and reference, the module generates a query field containing a concise summary of learned lessons, ensuring efficient retrieval and application of insights in subsequent decision-making endeavors.

Table 2: Differences between reflection of low and high

Reflection	Low-level Reflection	High-level Reflection
Target	Price Movements	Trading Decisions
Visual Data	Kline Chart	Trading Chart
Market Understanding	Micro	Macro
Function	Adaptability	Amendability

High-level Reflection Module The high-level reflection module is designed to provide analysis and reflections on past trading decisions. Besides the past trading decisions and their underlying reasoning, this module incorporates a graphical representation of buy and sell points on a trading chart, coupled with a cumulative return plot, to offer an intuitive representation of the efficacy of historical decisions. The initial phase assesses each trading decision’s correctness, identifying successes and mistakes. Subsequently, the module recommends improvements or corrective actions tailored to each identified mistake or success, fostering a continuous learning process. Beyond individual decision analysis, the module generates

overarching lessons from both successes and mistakes, providing a summary that can be adapted to future trading decisions and a query text to facilitate the retrieval of relevant reflections. This iterative learning process equips the agent with a dynamic knowledge base that evolves with each decision and allows the trading agent to draw connections between similar scenarios, applying learned lessons for more informed decision-making.

4.4 Tool-Augmented Decision-making Module

The decision-making module integrates key inputs, including market intelligence summaries, low-level reflection about price movement analyses, and reflections on past decisions. Augmented tools with professional investment guidance and traditional trading strategies like MACD Crossover, KDJ with RSI Filter and Mean Reversion are also considered. The module analyzes sentiment in market intelligence, predicts bullish or bearish trends from price movements, reflects on lessons learned, and evaluates professional guidance and traditional indicators. Decisions are derived from combining insights from these analyses, also considering the current financial position, leading to a final decision—whether to buy, sell, or hold the asset. Leveraging the Chain-of-Thought (COT) approach and in-context learning principles, our trading decision-making module not only executes trades but also provides reasoning, ensuring that each decision is rooted in a comprehensive understanding of market dynamics and contextual knowledge.

5 EXPERIMENT SETUP

Our research aims to conduct a thorough evaluation of FinAgent’s trading effectiveness, underscoring its unique capability to function efficiently with a significantly reduced historical data training window. This assessment also involves leveraging multimodal data inputs, incorporating both informational and agent-assistive augmented tools, along with a multi-perspective diversified retrieval. This approach is intended to enhance the understanding of market dynamics and sentiments, enabling more comprehensive and logical decision-making processes along with substantiated explanations. To validate its effectiveness, we have conducted a series of experiments to address the following research questions (RQs):

- **RQ1:** Is FinAgent outperforming current state-of-the-art trading agents and handling tasks that challenge other algorithms?
- **RQ2:** What is the effectiveness of each component of FinAgent in contributing to its overall performance?
- **RQ3:** Does the integration of augmented tools in FinAgent lead to a distinguishable improvement in its trading performance?
- **RQ4:** How effective is the diversified retrieval in FinAgent?

5.1 Datasets

Table 3: Dataset statistics detailing the chronological period and the number of each data source for each asset.

Asset	AAPL	AMZN	GOOGL	MSFT	TSLA	ETHUSD
Trading Date	From 2022-06-01 to 2024-01-01 (398 trading days)					
Asset Price	398 × (open, high, low, close, adj_close)					
Visual Data	398 × (Kline Chart, Trading Chart)					
Asset News	9748	10007	7923	8178	10076	2611
Expert Guidance	593	509	488	393	600	–

To conduct a thorough evaluation of FinAgent, we evaluate it across 6 real-world datasets. These included five datasets from the US stock markets, and one is the cryptocurrency. Each of them has multiple forms of data that come from various sources. Specifically, i) **Asset Price** at the day-level, including price data for open, high, low, close, and adj close. ii) **Visual Data** consists of historical Kline charts and trading charts, which are visual representations of asset market data and trading process on a daily basis. iii) **Asset News** coverage with daily updates from various esteemed sources such as Bloomberg Technology, Seeking Alpha and CNBC Television, ensuring a diverse and thorough perspective on the financial markets. iv) **Expert Guidance** provided by financial experts as the auxiliary information, aiming to furnish a thorough and well-rounded comprehension of market status. We summarize statistics of the 6 datasets in Table 3 and further elaborate on them in Appendix A.

Our diversified portfolio includes five major stocks: Apple Inc. (AAPL), Amazon.com Inc. (AMZN), Alphabet Inc. (GOOGL), Microsoft Corporation (MSFT), and Tesla Inc. (TSLA) and a prominent cryptocurrency named Ethereum (ETHUSD). This selection aims to showcase FinAgent’s versatility and consistency across various financial assets. Chosen for their extensive news coverage and representation of different market sectors, these data provide a robust basis for assessing FinAgent’s generalization capabilities across diverse financial environments. For dataset split, the data from the latter half of the year is allocated for testing (2023-06-01 ~ 2024-01-01) purposes, while the data from the penultimate year is utilized for training (2022-06-01 ~ 2023-06-01).

5.2 Evaluation Metrics

We compare FinAgent and baselines in terms of 6 financial metrics following [29, 38], which include 1 profit metric: annual return rate (ARR), 3 risk-adjusted profit metrics: Sharpe ratio (SR), Calmar ratio (CR), Sortino ratio (SoR), and 2 risk metrics: maximum drawdown (MDD), volatility (VOL). Definitions and formulas are as follows:

- **Annual Rate of Return (ARR)** is the annualized average return rate, calculated as $ARR = \frac{V_T - V_0}{V_0} \times \frac{C}{T}$, where T is the total number of trading days, and $C = 252$ is the number of trading days within a year. V_T and V_0 represent the final and initial portfolio values.
- **Sharpe Ratio (SR)** measures risk-adjusted returns of portfolios. It is defined as $SR = \frac{\mathbb{E}[\mathbf{r}]}{\sigma[\mathbf{r}]}$, where $\mathbb{E}[\cdot]$ is the expectation, $\sigma[\cdot]$ is the standard deviation, $\mathbf{r} = [\frac{V_1 - V_0}{V_0}, \frac{V_2 - V_1}{V_1}, \dots, \frac{V_T - V_{T-1}}{V_{T-1}}]^T$ denotes the historical sequence of the return rate.
- **Volatility (VOL)** is the variation in an investment’s return over time, measured as the standard deviation $\sigma[\mathbf{r}]$.
- **Maximum Drawdown (MDD)** measures the largest loss from any peak to show the worst case. It is defined as: $MDD = \max_{i=0}^T \frac{P_i - R_i}{P_i}$, where $R_i = \prod_{j=1}^i \frac{V_j}{V_{j-1}}$ and $P_i = \max_{k=1}^i R_k$.
- **Calmar Ratio (CR)** compares average annualized return to maximum drawdown, assessing risk-adjusted performance. It is defined as $CR = \frac{\mathbb{E}[\mathbf{r}]}{MDD}$.
- **Sortino Ratio (SoR)** is a risk-adjusted measure that focuses on the downside risk of a portfolio. It is defined as $SoR = \frac{\mathbb{E}[\mathbf{r}]}{DD}$, where DD is the standard deviation of negative return.


5.3 Baselines

We compare and evaluate the trading performance of FinAgent with four widely accepted conventional rule-based trading strategies (**B&H**, **MACD**, **KDJ&RSI** and **ZMR**) and eight advanced algorithms. Among these, price prediction models based on machine learning and deep learning (ML & DL-based) include **LGBM**[52], **LSTM**[52], and **Transformer**[52]. **SAC** [11], **PPO** [34] and **DQN** [21] are three models employed deep reinforcement learning (RL-based) methods, **FinGPT** [50] is based on LLM, and another is **FinMem** [56] that based on LLM Agents. The following will provide a brief introduction to each model:

- **Rule-based**
 - **Buy-and-Hold (B&H)** involves holding assets for an extended period, regardless of short-term market fluctuations, assuming that long-term returns will be more favorable.
 - **Moving Average Convergence Divergence (MACD)** is a technical analysis tool that uses MACD indicator and signal line crossovers to identify trading signals and market trends.
 - **KDJ with RSI Filter (KDJ&RSI)** integrates the KDJ indicator for detecting market extremes with the RSI indicator for momentum analysis to identify precise trading signals in financial markets.
 - **Z-score Mean Reversion (ZMR)** assumes that the price will revert to its mean over time with the metric of Z-score.
- **ML&DL-based**
 - **LGBM** [52] uses a series of tree models to predict price fluctuations and provide buy and sell signals.
 - **LSTM** [52] utilizes long short-term memory to improve the accuracy of price predictions.

Table 4: Performance comparison of all methods on six profitable metrics. Results in red, yellow and green show the best, second best and third best results on each dataset. The improvement row is the FinAgent over the best-performing baselines.

Categories	Models	AAPL			AMZN			GOOGL			MSFT			TSLA			ETHUSD		
		ARR% \uparrow	SR \uparrow	MDD% \downarrow	ARR% \uparrow	SR \uparrow	MDD% \downarrow	ARR% \uparrow	SR \uparrow	MDD% \downarrow	ARR% \uparrow	SR \uparrow	MDD% \downarrow	ARR% \uparrow	SR \uparrow	MDD% \downarrow	ARR% \uparrow	SR \uparrow	MDD% \downarrow
Market	B&H	13.0	0.6	14.78	42.33	1.08	17.38	22.47	0.71	12.97	22.49	0.84	12.92	37.4	0.72	32.65	29.26	0.87	23.21
Rule-based	MACD	11.86	0.72	10.38	14.27	0.71	7.84	-18.0	-0.89	20.07	15.23	0.77	8.34	-4.9	-0.02	14.15	10.24	0.47	24.32
	KDJ&RSI	2.17	0.17	11.88	19.38	0.65	17.27	24.39	2.13	2.03	18.84	1.06	7.78	2.14	0.17	24.73	8.87	0.51	16.95
	ZMR	-3.91	-0.22	8.88	18.73	0.84	7.89	32.51	1.45	5.38	9.86	0.71	6.22	-7.28	-0.09	19.9	29.35	1.23	13.11
ML&DL-based	LGBM	16.93	1.47	2.52	29.34	0.72	17.41	24.77	0.7	12.98	19.28	0.67	12.96	15.57	0.84	3.88	24.91	0.72	22.96
	LSTM	10.97	0.54	11.95	15.91	0.46	17.41	24.86	0.7	12.98	18.86	0.68	11.75	17.36	0.78	4.44	36.09	1.03	21.5
	Transformer	17.11	0.96	7.53	32.66	1.11	4.96	13.69	0.46	12.93	17.44	1.46	2.59	39.7	1.04	8.17	31.0	1.02	12.93
RL-based	DQN	7.92	0.4	14.88	27.43	1.17	5.27	34.4	1.39	7.15	30.44	1.18	10.56	15.07	0.44	28.12	29.81	1.18	9.53
	SAC	24.84	1.12	11.98	38.33	1.07	13.84	23.8	0.75	13.07	22.02	0.82	12.92	42.22	0.87	26.19	17.84	0.76	10.06
	PPO	13.26	0.61	14.78	21.17	0.7	13.84	38.29	1.3	8.45	11.32	0.48	17.51	33.64	0.78	28.35	34.75	1.31	11.12
LLM-based	FinGPT	-5.46	-0.17	16.23	42.93	1.1	18.94	12.28	0.44	13.0	25.1	0.97	9.84	38.43	0.75	31.47	21.57	0.68	25.56
	FinMem	23.78	1.11	10.39	40.07	1.03	18.53	31.27	1.11	8.97	40.58	1.5	7.48	50.04	0.92	25.77	44.72	1.27	13.59
Ours	FinAgent	31.9	1.43	10.4	65.1	1.61	13.2	56.15	1.78	8.45	44.74	1.79	5.57	92.27	2.01	12.14	43.08	1.18	12.72
Improvement(%)		28.39	-	-	51.64	37.61	-	46.64	-	-	10.25	19.33	-	84.39	93.27	-	-	-	-

**Figure 4: Performance comparison over time between FinAgent and other benchmarks across all assets.**

- **Transformer** [52] models leverage self-attention mechanisms to enhance the precision of price forecasts.
- **RL-based**
 - **SAC** [11] is an off-policy actor-critic algorithm that optimizes trading strategies using entropy regularization and soft value functions in continuous action spaces.
 - **PPO** [34] updates trading policies iteratively to balance exploration and exploitation, ensuring stability and sample efficiency.
 - **DQN** [21] uses deep neural networks to approximate the action-value function and make trading decisions from market data.
- **LLM-based**
 - **FinGPT** [50] is an open-source LLM framework for converting financial news and prices into financial decisions.
 - **FinMem** [56] is an advanced LLM agent framework for automated trading, fine-tuned to boost investment returns.

5.4 Implementation Details

Although FinAgent’s training and inference can be done without a GPU, we utilized a single NVIDIA RTX A6000 GPU for our benchmark methods. To ensure equitable comparison, all benchmarks are conducted within the same RL environment for both training and evaluation. The following experiments related to FinAgent all have diversified retrieval if not specifically noted. Details on the benchmark and experiments setup are provided in Appendix C.

6 EXPERIMENTAL RESULTS

Comparison with Baselines (RQ1). We compared FinAgent with 9 baseline methods in terms of 6 financial metrics. Table 4 and Figure 4 demonstrate our method significantly outperforms existing baselines, especially remarkable improvements in profitability, and setting a new benchmark in the field. The full results and case studies of FinAgent are available in Appendix B. FinAgent’s performance on the five stocks, as measured by ARR% and SR, with enhancements of at least 10% and 19%, compared to the best-performing baseline, respectively. Notably, its performance on the TSLA dataset stands out even more, achieving 84% and 118% improvement, significantly outperforming all other baselines. Across all datasets, FinAgent is the only method that consistently outperforms the broader market in terms of profitability. In contrast, FinMem falls short on the AMZN dataset, where its ARR% is 40%, underperforming the market’s Buy & Hold (B&H) strategy at 42%. This underscores the superior stability and robustness of FinAgent compared to other baselines. We can also observe that rule-based methods are optimal in controlling risk, but not outstanding in capturing returns. This is because rule-based model methods are robust to outliers and noise in the data and thus can reduce decision risk. It is worth noting that high returns often come with high risks. Hence, FinAgent represents a slight compromise on risk control. This result relates to our chosen investor preference of an aggressive trader. Therefore,

FinAgent can take on slightly higher risk to achieve substantially greater returns. It allows FinAgent to optimize performance by balancing risk and reward effectively.

Figure 4 illustrates that FinAgent’s performance surpasses other methods regarding cumulative returns, particularly on the TSLA dataset. Leveraging market intelligence and the reflection mechanism, FinAgent anticipates a significant stock price drop post-September 14, 2023. By taking a short position, it can effectively hedge against potential trading losses and generate high returns.

It’s important to note that our approach yields slightly lower returns than FinMem on the cryptocurrency ETH, primarily because our auxiliary agents are specialized strategies tailored for stocks, not for cryptocurrencies with higher trading frequency. Further insights from the ablation study section for FinAgent reveal that employing a generalized auxiliary agent for cryptocurrency could potentially increase returns to 54%, compared to the current 44%. This significant difference will be elaborated upon in the forthcoming ablation studies.

7 ABLATION STUDIES

7.1 Effectiveness of Each Component (RQ2)

In Table 5, we study the effectiveness of market intelligence (M), low-level reflection (L), high-level reflection (H) and augmented tools (T). When compared to using solely M and ML , the integration of the low-level reflection module leads to an impressive increase in ARR% by 45% to 101% for TSLA, and ETHUSD, and cutting risk by 14% to 44%. When comparing the ML and MLH , the addition of the high-level reflection module significantly enhances the ARR% and SR, while notably reducing risk. This improvement comes with a minor trade-off: a slight 7% rise in MDD% for TSLA. Compared to MLH and $MLHT$, there’s a minor improvement in stock profitability. However, the performance of ETH cryptocurrency dropped by over 20% due to the introduction of rule-based methods as auxiliary agents, which are specialized only for stocks.

Table 5: Ablation studies over different components. \checkmark indicates adding the component to FinAgent. Red and green indicate performance improvement and reduction.

M L H T	TSLA			ETHUSD		
	ARR% \uparrow	SR \uparrow	MDD% \downarrow	ARR% \uparrow	SR \uparrow	MDD% \downarrow
\checkmark	39.01	0.90	22.54	16.21	0.63	15.93
\checkmark	39.27	0.77	30.15	25.97	0.77	24.43
$\checkmark\checkmark$	57.16(+45.56%)	1.02(+33.14%)	25.77(-14.52%)	52.33(+101.48%)	1.34(+72.99%)	13.59(-44.39%)
$\checkmark\checkmark\checkmark$	89.25(+56.14%)	1.46(+42.86%)	27.62(+7.18%)	54.80(+4.73%)	1.40(+5.09%)	11.74(-13.57%)
$\checkmark\checkmark\checkmark\checkmark$	92.27(+3.38%)	2.01(+37.84%)	12.14(-56.04%)	43.08(-21.39%)	1.18(-16.09%)	12.72(+8.30%)

7.2 Effectiveness of Augmented Tools (RQ3)

As previously discussed, while the addition of auxiliary agents to stock investments results in profit improvements, it causes a considerable performance decline in cryptocurrencies. Thus, we conduct the experiment that decisions are made solely by augmented tools, such as rule-based methods serving as auxiliary agents. We conducted the experiment in which various auxiliary agents provided both decisions and their explanations. These inputs are directly integrated into FinAgent’s decision-making module without other modules’ involvement in the final decision process. As shown in Table 4 and Table 5, the 16% ARR% for solely T method starkly

contrasts with the 29% ARR% of B&H in ETHUSD, highlighting the inefficacy of the stock-specific rule-based methods for cryptocurrencies and demonstrating that introducing to FinAgent significantly affects performance. This suggests that investors should not indiscriminately add auxiliary agents for investment support. Instead, they must meticulously select agents that match the characteristics of the market to avoid detrimental impact on performance.

7.3 Effectiveness of Diversified Retrieval (RQ4)

As shown in Figure 5(a), we compare the performance of FinAgent with or without diversified retrieval on AAPL, and find that the use of diversified retrieval can contribute an obvious improvement in ARR and SR. As shown in Figure 5(b), we extract different types of market intelligence that AAPL diversified retrieve to daily on the validation set and filter out individuals with the same content under the same type. We perform t-SNE visualization of its LLM extracted embedding, and we can find that the LLM extracted embedding has a clear distinction between different retrieval types, which proves the effectiveness of our method.




Figure 5: (a) Performance of FinAgent with/without diversified retrieval on AAPL. (b) Visualization of diversified retrieved market intelligence embedding by t-SNE on AAPL.

8 CONCLUSION AND FUTURE WORK

This paper introduces FinAgent, a financial trading agent powered by LLM that exhibits high reasoning ability and generalizability. FinAgent is a multimodal agent that integrates both textual and visual data, enabling a comprehensive understanding of market dynamics and historical trading behaviors. It is designed to independently leverage auxiliary tools for detailed market data analysis over different time scales. With its multi-perspective and diverse retrieval approach, FinAgent effectively identifies correlations between current market conditions and past market patterns and trends and integrates market information to make final and effective decisions. For future research directions, we will apply FinAgent to other financial tasks, such as portfolio management, where LLM is used to rank each stock according to the observed market intelligence and make the stock selection.

9 ACKNOWLEDGMENTS

This project is supported by the National Research Foundation, Singapore under its Industry Alignment Fund – Pre-positioning (IAF-PP) Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A Next-Generation Hyperparameter Optimization Framework. In *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2623–2631.
- [2] Bo An, Shuo Sun, and Rundong Wang. 2022. Deep reinforcement learning for quantitative trading: Challenges and opportunities. *IEEE Intelligent Systems* 37, 2 (2022), 23–26.
- [3] Chi Chen, Li Zhao, Jiang Bian, Chunxiao Xing, and Tie-Yan Liu. 2019. Investment behaviors can tell what inside: Exploring stock intrinsic properties for stock trend prediction. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2376–2384.
- [4] Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. 2023. Theoremqa: A theorem-driven question answering dataset. *arXiv preprint arXiv:2305.12524* (2023).
- [5] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* 24, 240 (2023), 1–113.
- [6] Filippos Christianos, Georgios Papoudakis, Matthieu Zimmer, Thomas Coste, Zhihao Wu, Jingxuan Chen, Khyati Khandelwal, James Doran, Xidong Feng, Jiacheng Liu, Zheng Xiong, Yicheng Luo, Jianye Hao, Kun Shao, Haitham Bou-Ammar, and Jun Wang. 2023. Pangu-Agent: A fine-tunable generalist agent with structured reasoning. *arXiv:2312.14878* [cs.AI]
- [7] Yue Deng, Feng Bao, Youyong Kong, Zhiquan Ren, and Qionghai Dai. 2016. Deep direct reinforcement learning for financial signal representation and trading. *IEEE Transactions on Neural Networks and Learning Systems* 28, 3 (2016), 653–664.
- [8] Yi Ding, Weiqing Liu, Jiang Bian, Daoqiang Zhang, and Tie-Yan Liu. 2018. Investor-imitator: A framework for trading knowledge extraction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1310–1319.
- [9] Robert D Edwards, John Magee, and WH Charles Bassetti. 2018. *Technical analysis of stock trends*. CRC press.
- [10] Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. 2019. Temporal relational ranking for stock prediction. *ACM Transactions on Information Systems (TOIS)* 37, 2 (2019), 1–30.
- [11] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
- [12] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiaowu Zheng, Yuheng Cheng, Ceyao Zhang, Jinlin Wang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2023. MetaGPT: Meta programming for a multi-agent collaborative framework. *arXiv:2308.00352* [cs.AI]
- [13] Ziniu Hu, Weiqing Liu, Jiang Bian, Xuanzhe Liu, and Tie-Yan Liu. 2018. Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction. In *Proceedings of the 11th ACM International Conference on Web Search and Data Mining (WSDM)*. 261–269.
- [14] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems* 30 (2017).
- [15] Xiao-Yang Liu, Hongyang Yang, Qian Chen, Runjia Zhang, Liuqing Yang, Bowen Xiao, and Christina Dan Wang. 2020. FinRL: A deep reinforcement learning library for automated stock trading in quantitative finance. *Deep RL Workshop, NeurIPS 2020* (2020).
- [16] Yang Liu, Qi Liu, Hongke Zhao, Zhen Pan, and Chuanren Liu. 2020. Adaptive quantitative trading: An imitative deep reinforcement learning approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 2128–2135.
- [17] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2023. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255* (2023).
- [18] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems* 35 (2022), 2507–2521.
- [19] Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. Chameleon: Plug-and-play compositional reasoning with large language models. *arXiv:2304.09842* [cs.CL]
- [20] Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. 2022. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. *arXiv preprint arXiv:2209.14610* (2022).
- [21] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [22] Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332* (2021).
- [23] David MQ Nelson, Adriano CM Pereira, and Renato A de Oliveira. 2017. Stock market’s price movement prediction with LSTM neural networks. In *2017 International Joint Conference on Neural Networks (IJCNN)*. 1419–1426.
- [24] OpenAI. 2021. Chatgpt. <https://openai.com/research/chatgpt>
- [25] OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774* [cs.AI]
- [26] OpenAI. 2023. GPT-4V(ision) system card. <https://openai.com/research/gpt-4v-system-card>
- [27] Aaron Parisi, Yao Zhao, and Noah Fiedel. 2022. Talm: Tool augmented language models. *arXiv preprint arXiv:2205.12255* (2022).
- [28] Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive simulacra of human behavior. *arXiv:2304.03442* [cs.HC]
- [29] Molei Qin, Shuo Sun, Wentao Zhang, Haochong Xia, Xinrun Wang, and Bo An. 2023. Earnhft: Efficient hierarchical reinforcement learning for high frequency trading. *arXiv preprint arXiv:2309.12891* (2023).
- [30] Ramit Sawhney, Shivam Agarwal, Arnab Wadhwa, and Rajiv Shah. 2020. Deep attentive learning for stock movement prediction from social media text and company correlations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 8415–8426.
- [31] Ramit Sawhney, Piyush Khanna, Arshiya Aggarwal, Taru Jain, Puneet Mathur, and Rajiv Shah. 2020. VolTAGE: Volatility forecasting via text-audio fusion with graph convolution networks for earnings calls. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 8001–8013.
- [32] Ramit Sawhney, Arnab Wadhwa, Shivam Agarwal, and Rajiv Shah. 2021. Quantitative day trading from natural language using reinforcement learning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 4018–4030.
- [33] Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761* (2023).
- [34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [35] Theodore R Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L Griffiths. 2023. Cognitive architectures for language agents. *arXiv preprint arXiv:2309.02427* (2023).
- [36] Liangtai Sun, Yang Han, Zihan Zhao, Da Ma, Zhennan Shen, Baocai Chen, Lu Chen, and Kai Yu. 2023. Scieval: A multi-level large language model evaluation benchmark for scientific research. *arXiv preprint arXiv:2308.13149* (2023).
- [37] Shuo Sun, Molei Qin, Xinrun Wang, and Bo An. 2023. PRUDEX-Compass: Towards systematic evaluation of reinforcement learning in financial markets. *Transactions on Machine Learning Research* (2023).
- [38] Shuo Sun, Molei Qin, wentao zhang, Haochong Xia, Chuqiao Zong, Jie Ying, Yonggang Xie, Lingxuan Zhao, Xinrun Wang, and Bo An. 2023. TradeMaster: A holistic quantitative trading platform empowered by reinforcement learning. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [39] Shuo Sun, Rundong Wang, and Bo An. 2023. Reinforcement learning for quantitative trading. *ACM Transactions on Intelligent Systems and Technology* 14, 3 (2023), 1–29.
- [40] Shuo Sun, Xinrun Wang, Wanqi Xue, Xiaoxuan Lou, and Bo An. 2023. Mastering stock markets with efficient mixture of diversified trading experts. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*. 2109–2119.
- [41] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulkshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. 2022. Llama: Language models for dialog applications. *arXiv preprint arXiv:2201.08239* (2022).
- [42] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [43] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv: Arxiv-2305.16291* (2023).
- [44] Jia Wang, Tong Sun, Benyuan Liu, Yu Cao, and Hongwei Zhu. 2019. CLVSA: A convolutional LSTM based variational sequence-to-sequence model with attention for predicting trends of financial markets. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 3705–3711.
- [45] Rundong Wang, Hongxin Wei, Bo An, Zhouyan Feng, and Jun Yao. 2021. Commission fee is not enough: A hierarchical reinforced framework for portfolio management. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 626–633.

- [46] Xiaoxuan Wang, Ziniu Hu, Pan Lu, Yanqiao Zhu, Jieyu Zhang, Satyen Subramaniam, Arjun R Loomba, Shichang Zhang, Yizhou Sun, and Wei Wang. 2023. Scibench: Evaluating college-level scientific problem-solving abilities of large language models. *arXiv preprint arXiv:2307.10635* (2023).
- [47] Zhicheng Wang, Biwei Huang, Shikui Tu, Kun Zhang, and Lei Xu. 2021. Deep-Trader: a deep reinforcement learning approach for risk-return balanced portfolio management with market conditions Embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 643–650.
- [48] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. 2023. Bloomberggpt: A large language model for finance. *arXiv preprint arXiv:2303.17564* (2023).
- [49] Yumo Xu and Shay B Cohen. 2018. Stock movement prediction from tweets and historical prices. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*. 1970–1979.
- [50] Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023. FinGPT: Open-Source Financial Large Language Models. *arXiv preprint arXiv:2306.06031* (2023).
- [51] Hui Yang, Sifu Yue, and Yunzhong He. 2023. Auto-GPT for Online Decision Making: Benchmarks and Additional Opinions. *arXiv:2306.02224 [cs.AI]*
- [52] Xiao Yang, Weiqing Liu, Dong Zhou, Jiang Bian, and Tie-Yan Liu. 2020. Qlib: An AI-oriented quantitative investment platform. *arXiv preprint arXiv:2009.11189* (2020).
- [53] Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023. MM-REACT: Prompting ChatGPT for multimodal reasoning and action. *arXiv:2303.11381 [cs.CV]*
- [54] Yunan Ye, Hengzhi Pei, Boxin Wang, Pin-Yu Chen, Yada Zhu, Ju Xiao, and Bo Li. 2020. Reinforcement-learning based portfolio management with augmented asset movement prediction states. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 1112–1119.
- [55] Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2022. Generate rather than retrieve: Large language models are strong context generators. *arXiv preprint arXiv:2209.10063* (2022).
- [56] Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W. Suchow, and Khaldoun Khashanah. 2023. FinMem: A performance-enhanced LLM trading agent with layered memory and character design. *arXiv:2311.13743 [q-fin.CP]*
- [57] Chi Zhang, Zhao Yang, Jiakuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2023. AppAgent: Multimodal agents as smartphone users. *arXiv:2312.13771 [cs.CV]*

A DETAILS OF DATASETS AND PROCESSING

To conduct a thorough evaluation of FinAgent, we evaluate it across 6 real-world datasets. These included five datasets from the US stock markets and one is the cryptocurrency. Each of them have multiple forms of data that come from various sources. Specifically, i) **Asset Price** at the day-level, including price data for open, high, low, close, and adj close; ii) **Visual Data** consists of historical Kline charts and trading charts, which are visual representations of asset market data and trading process on a daily basis; iii) **Asset News** coverage with daily updates from various esteemed sources, including Bloomberg Technology, Seeking Alpha, CNBC Television, and more, ensuring a diverse and thorough perspective on the financial markets; iv) **Expert Guidance** provided by financial experts as the auxiliary information, aiming to furnish a thorough and well-rounded comprehension of market status. We summarize statistics of the 6 datasets in Table 3 and further elaborate on them as follows:

Asset. We selected a varied portfolio comprising five stocks Apple Inc. (AAPL), Amazon.com Inc. (AMZN), Alphabet Inc. (GOOGL), Microsoft Corporation (MSFT), and Tesla Inc. (TSLA), a foreign exchange pair, and a prominent cryptocurrency, Ethereum (ETH). This selection aims to showcase FinAgent’s versatility and consistency across various financial assets. Chosen for their extensive news coverage and representation of different market sectors, these data provide a robust basis for assessing FinAgent’s generalization capabilities across diverse financial environments.

Price and News. We acquired price and news data for all assets from Financial Modeling Prep⁵ (FMP), wherein the price data encompasses including open, high, low, close, and adj close. The news data is sourced from renowned market analysis and stock research platforms, notably including Seeking Alpha and so on. This selection ensures a comprehensive dataset, integrating both quantitative financial metrics and qualitative market insights.

Visual Data. Within the textual data framework, we furnish FinAgent with visual information, specifically including historical Kline charts and trading line charts, to enhance its analytical capabilities. The tool employed for this plotting task is the pyecharts⁶, a specialized library for financial data visualization.

Expert Guidance. Expert Guidance is provided as a distinct component of the auxiliary information by augmented tools. This selection ensures a comprehensive dataset, integrating professional analysts and individual investors insights, fostering a diverse range of perspectives in the investment community. We obtained the expert professional analysis from Seeking Alpha⁷. Seeking Alpha is a popular platform among investors and financial analysts, is renowned for its diverse professional analysis, providing valuable insights from seasoned analysts across the financial market spectrum.

Trading Date. For dataset split, the data from the latter half of the year is allocated for testing (2023-06-01⁸ ~ 2024-01-01) purposes, while the data from the penultimate year is utilized for training (2022-06-01 ~ 2023-06-01).

B DETAILS OF COMPARISON WITH BASELINES

We compared FinAgent with 9 baseline methods in terms of 6 financial metrics. Table 6 and Figure 4 demonstrate our method significantly outperforms existing baselines, especially remarkable improvements in profitability, and setting a new benchmark in the field.

C DETAILS OF IMPLEMENTATION

Although FinAgent’s training and inference can be done without a GPU, we utilized a single NVIDIA RTX A6000 GPU for our benchmark methods. For dataset split, the data from the latter half of the year is allocated for testing (2023-06-01 ~ 2024-01-01) purposes, while the data from the penultimate year is utilized for training (2022-06-01 ~ 2023-06-01). To ensure equitable comparison, all benchmarks are conducted within the same RL environment for both training and evaluation.

Benchmark Setup. In the training phase, we use OPTUNA [1] for hyperparameter optimization, adapting both rule-based and RL methods to trading. This process is followed by an evaluation with the optimal parameters. We employ the officially provided default parameters for both training and testing of FinGPT and FinMem.

FinAgent Setup. For each training dataset, we perform only one round of training without the usual requirement for multiple iterations in fine-tuning LLMs. As demonstrated by FinMem [56],

⁵FMP API provides data about stock historical price and news, company financial statements, and cryptocurrencies. Entry is <https://site.financialmodelingprep.com>.

⁶<https://github.com/pyecharts/pyecharts>

⁷<https://seekingalpha.com/>

⁸Dates follow the YYYY-MM-DD format, e.g., "2023-06-01" for June 1st, 2023.

Table 6: Performance comparison of all methods on six profitable metrics. Results in red, yellow and green show the best, second best and third best results on each dataset. Improvement is the FinAgent over the best-performing baselines.

Categories	Models	AAPL			AMZN			GOOGL			MSFT			TSLA			ETHUSD		
		ARR%↑	SR↑	MDD%↓	ARR%↑	SR↑	MDD%↓	ARR%↑	SR↑	MDD%↓	ARR%↑	SR↑	MDD%↓	ARR%↑	SR↑	MDD%↓	ARR%↑	SR↑	MDD%↓
Market	B&H	13.0024	0.5998	14.7809	42.3337	1.0834	17.3848	22.4726	0.7108	12.9705	22.4942	0.8373	12.9214	37.4009	0.7239	32.6523	29.2588	0.8655	23.2077
Rule-based	MACD	11.8642	0.7221	10.3799	14.2748	0.7056	7.841	-18.0034	-0.8867	20.0718	15.2322	0.7704	8.3445	-4.8974	-0.0203	14.1546	10.236	0.4689	24.3238
	KDJ&RSI	2.1737	0.1746	11.8789	19.3757	0.6495	17.2746	24.391	2.1282	2.03	18.8415	1.0587	7.7806	2.137	0.1695	24.727	8.8745	0.5098	16.9536
	ZMR	-3.9084	-0.2186	8.8819	18.7289	0.8412	7.8938	32.5112	1.4533	5.3845	9.8637	0.7106	6.221	-7.2806	-0.0863	19.9048	29.3519	1.2294	13.1098
ML&DL-based	LGBM	16.9268	1.4708	2.5204	29.3395	0.7187	17.414	24.7746	0.6958	12.9814	19.2771	0.6668	12.9616	15.575	0.843	3.8844	24.9111	0.7154	22.9568
	LSTM	10.9742	0.5363	11.9535	15.9051	0.4588	17.414	24.8583	0.6989	12.9814	18.8603	0.6779	11.7544	17.3617	0.7796	4.4384	36.0865	1.0254	21.5043
	Transformer	17.115	0.957	7.5295	32.6621	1.1134	4.9593	13.692	0.4571	12.9253	17.4417	1.4553	2.5895	39.701	1.0445	8.1721	31.0038	1.0205	12.9309
RL-based	DQN	7.9236	0.401	14.8785	27.4305	1.1701	5.2736	34.4026	1.3859	7.1473	30.4406	1.1782	10.5612	15.0693	0.443	28.1204	29.8052	1.1826	9.5297
	SAC	24.8449	1.1234	11.9776	38.3318	1.0733	13.8432	23.8034	0.7506	13.0667	22.0218	0.8177	12.9214	42.2209	0.8727	26.1947	17.8439	0.7635	10.0587
	PPO	13.2619	0.6096	14.7809	21.1745	0.6965	13.8432	38.2907	1.2982	8.4536	11.3219	0.4831	17.5054	33.6444	0.7767	28.3527	34.7469	1.3096	11.1171
LLM-based	FinGPT	-5.4632	-0.1731	16.2268	42.9331	1.1026	18.9359	12.277	0.4444	13.0013	25.1012	0.9667	9.8426	38.4338	0.7504	31.474	21.5746	0.6801	25.562
	FinMem	23.7809	1.1073	10.3872	40.07	1.034	18.5279	31.2716	1.1073	8.9706	40.5757	1.4989	7.4838	50.0353	0.9233	25.7714	44.717	1.2738	13.587
FinAgent (Our)	No-finetuned	-2.0047	-0.0216	16.7567	41.6548	1.3959	5.2678	21.108	0.682	13.0719	22.6483	0.8373	12.9214	38.0164	1.2951	12.0609	23.4038	0.7924	23.231
	w/o-MLH	18.5186	0.9882	11.8842	62.3106	1.6379	11.3999	37.3308	1.2657	4.0	18.297	0.8945	9.6792	39.0087	0.8971	22.5376	16.2106	0.6329	15.9325
	w/o-LHT	12.6872	0.5849	14.7806	43.2195	1.1057	14.6786	17.414	0.5763	13.006	18.4571	0.7108	14.2551	39.2704	0.7662	30.1507	25.9708	0.7718	24.4514
	w/o-HT	21.3044	0.9777	13.6424	47.3916	1.2032	15.5038	29.3235	1.0177	8.9706	39.9073	1.4777	7.4838	57.1638	1.0201	25.7714	52.3265	1.3351	13.587
	w/o-T	33.7509	1.5205	8.1783	63.8116	1.7009	9.7325	52.1066	1.8228	8.3106	42.5213	1.4909	7.6998	89.2532	1.4573	27.6213	54.804	1.403	11.7427
FinAgent	31.8972	1.4326	10.4032	65.0998	1.6096	13.198	56.1542	1.7786	8.4532	44.7359	1.7884	5.5732	92.2677	2.0088	12.143	43.0822	1.1733	12.7171	
Improvement(%)		35.8464	3.3791	-	51.6308	45.3636	-	46.6523	-	-	10.2529	19.3142	-	84.4052	92.3217	-	22.5574	7.1739	-

Categories	Models	AAPL			AMZN			GOOGL			MSFT			TSLA			ETHUSD		
		SOR↑	CR↑	VOL↓	SOR↑	CR↑	VOL↓	SOR↑	CR↑	VOL↓	SOR↑	CR↑	VOL↓	SOR↑	CR↑	VOL↓	SOR↑	CR↑	VOL↓
Market	B&H	16.5846	0.9589	0.0114	35.1804	2.4319	0.0188	18.5186	1.9025	0.0167	26.5133	1.8135	0.0135	23.3319	1.3856	0.0301	23.2235	1.3831	0.0222
Rule-based	MACD	13.7755	1.1877	0.0082	19.2452	1.9176	0.0103	-20.917	-0.8883	0.0097	18.7485	1.8966	0.0099	-0.4067	-0.0544	0.0182	10.8951	0.5433	0.0166
	KDJ&RSI	3.3994	0.2578	0.0084	16.6915	1.2471	0.016	36.6655	11.4076	0.0052	19.5346	2.4169	0.0085	3.3617	0.2737	0.0192	7.7786	0.6106	0.012
	ZMR	-2.9977	-0.37	0.0072	9.5853	2.4473	0.011	35.1125	5.7782	0.0103	12.9275	1.6385	0.0069	-1.8214	-0.1642	0.0182	21.4433	2.1715	0.0137
ML&DL-based	LGBM	45.2444	6.5825	0.0049	26.2943	1.8469	0.0193	19.7497	2.0882	0.0167	22.6593	1.5993	0.0134	12.1981	4.0837	0.0081	21.8037	1.2503	0.0217
	LSTM	14.5498	0.9915	0.0095	14.7442	1.1241	0.0184	19.5489	2.0865	0.0167	21.5244	1.7144	0.0128	16.1653	4.0391	0.0099	27.1716	1.7031	0.0193
	Transformer	28.3604	2.2858	0.0078	27.8371	6.5258	0.0125	11.5053	1.2619	0.0154	19.7493	6.6012	0.0051	34.1884	4.8624	0.0164	27.864	2.4209	0.0166
RL-based	DQN	10.3705	0.6266	0.0111	29.7698	5.1156	0.0111	37.7389	4.619	0.0114	32.8642	2.8493	0.0122	11.0394	0.7399	0.0225	24.6279	3.0545	0.0143
	SAC	33.5676	2.0552	0.0105	32.432	2.766	0.0112	19.3937	1.9708	0.0165	26.0662	1.7829	0.0135	26.2869	1.441	0.0251	13.3918	1.9199	0.0146
	PPO	16.766	0.9747	0.0113	20.3167	1.6735	0.0159	42.7201	4.3691	0.0136	14.7808	0.7484	0.013	19.1547	1.3203	0.0231	27.2416	2.9836	0.0147
LLM-based	FinGPT	-4.6731	-0.246	0.0111	34.8082	2.2545	0.0186	11.6596	1.1842	0.0167	30.1935	2.5867	0.0127	23.1813	1.4402	0.0291	18.5817	1.0131	0.0221
	FinMem	29.8819	2.2731	0.0102	33.0779	2.1843	0.0188	34.7826	3.4572	0.0134	47.1061	5.1266	0.0123	25.8819	2.0887	0.028	34.1492	3.1349	0.0194
FinAgent (Our)	No-finetuned	-0.5635	-0.0295	0.011	46.9773	7.5705	0.0127	17.5637	1.7864	0.0164	26.6949	1.8259	0.0135	17.2111	3.0426	0.0136	19.6504	1.1081	0.0188
	w/o-MLH	20.8183	1.5678	0.009	53.4328	4.9724	0.0166	24.9489	2.9016	0.0137	22.0401	1.9273	0.01	20.3649	1.8355	0.0221	14.0758	1.1917	0.0174
	w/o-LHT	16.15	0.9396	0.0114	35.1674	2.9256	0.0186	15.1339	1.552	0.0168	22.4806	1.3883	0.0153	22.2082	1.5162	0.0286	21.0778	1.2131	0.0222
	w/o-HT	29.5194	1.576	0.0105	40.4896	2.9758	0.0184	33.6191	2.906	0.0139	46.4396	5.054	0.0123	28.3015	2.2956	0.0278	37.172	3.6027	0.0212
	w/o-T	46.6145	3.9301	0.0101	50.1096	5.9255	0.0162	62.2508	5.7201	0.0125	40.3937	5.2137	0.0129	41.5642	2.9906	0.0266	37.4619	4.2958	0.0208
FinAgent	44.2812	2.9424	0.0102	52.5602	4.8441	0.0176	62.2992	6.0365	0.0138	49.6249	7.4209	0.0111	45.4139	6.4543	0.0187	31.0159	3.3078	0.0207	
Improvement(%)		3.0282	-	-	51.8823	16.0088	-	45.8311	-	-	5.3471	12.4174	-	32.8342	32.7390	-	9.7007	37.0315	-

Table 7: Notations in the paper.

Notation	Description
t	Current day
T	Total trading days
$t - s, t + s$	Short-term price analysis from $t - s$ to t and t to s
$t - m, t + m$	Medium-term price analysis from $t - m$ to t and t to m
$t - l, t + l$	Long-term price analysis from $t - l$ to t and t to l
S	A finite set of states
s_t	State of day t
\mathcal{A}	A finite set of actions
a_t	Action of day t
\mathcal{T}	Transition function
R	Reward function
r_t	Reward of day t with s_t and a_t
γ	Discount factor
π	Policy
$\mu(\cdot)$	Specialized modules for reasoning
μ_t	Specialized modules of day t
λ	Financial trading task
Mem_t^λ	Memory of day t in the task λ
$Tool_t^\lambda$	Tool of day t in the task λ
$\phi(\cdot)$	Task-relevant prompt generator
\mathcal{D}_t^λ	Action parsing function
$M_t^\lambda, L_t^\lambda, H_t^\lambda$	M, L, H modules
$\phi_M^\lambda, \phi_L^\lambda, \phi_H^\lambda$	Prompt generator for M, L, H
$Mem_{M,t}^{\lambda}, Mem_{L,t}^{\lambda}, Mem_{H,t}^{\lambda}$	Memory of M, L, H modules of day t in the task λ
KC_t	Kline chart of day t
TC_t	Trading chart of day t
$SLMI_t$	Summary of latest market intelligence of day t
$QLMI_t = \{Q_1^L, \dots, Q_M^L\}$	M query texts for retrieving past market intelligence of day t
K	Retrieved topk items
$Q_{i,j}^P$	Retrieval type i and top j retrieved past market latest intelligence
$SPMI_t$	Summary of past market intelligence of day t
$LLR_t^{ST}, LLR_t^{MT}, LLR_t^{LT}$	Low-level reflection results at short term, medium term and long term impact
$QLLR_t$	Query text for low-level reflection of day t
$PLLR_t^{ST}, PLLR_t^{MT}, PLLR_t^{LT}$	retrieved topk low-level reflection in short term, medium term and long term
$HLLR_t$	High-level reflection results of day t
$QHLLR_t$	Query text for high-level reflection of day t
$PHLLR_t$	Retrieved topk high-level reflection of day t

OpenAI's GPT-4 shows improved performance over GPT-3.5. Consequently, we have selected GPT-4 as the foundational LLM for FinAgent. For the market intelligence and decision-making modules, which do not process visual data, we use the gpt-4-1106-preview. In contrast, the two reflection modules, which require an in-depth understanding of visual data, utilize gpt-4-vision-preview. For the memory module, which is designed to store and retrieve texts based on text similarity, a text encoder is essential for vectorizing the text. We adopt text-embedding-3-large for this purpose. The top-k of our retrieval samples is 5. In the low-level reflection module, short term, medium term and long term are for the latest 1 day, 7 days and 14 days respectively. It is important to note that past and future asset price increases and decreases are visible during the training phase, but only past trends are visible during the testing phase to prevent data leakage. The following experiments related to FinAgent all have diversified retrieval if not specifically noted.

D DETAILS OF NOTATIONS

We provide the main notations in Table 7.