Transactions on Computational and Scientific Methods | Vo. 5, No. 09, 2025
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Adaptive Stock Trading Algorithms Jointly Driven by Multi-Agent Collaborative Decision Making and Large Language Models in Complex Financial Market Environments

Shihao Sun
New York University, Brooklyn, USA
ss11972@nyu.edu

Abstract: To address the challenges of heterogeneous information sources, rapidly changing market conditions, and insufficient utilization of semantic information in traditional trading methods within complex financial markets, this paper proposes a multi-agent collaborative decision-making approach for automated stock trading, jointly driven by a large language model. This method, centered on the unified modeling requirements of structured market data and unstructured financial text, constructs an overall framework comprised of unified state encoding, agent interaction and collaboration, adaptive coordination and fusion, and action selection. This enables the effective fusion of multi-source content, including market prices, transaction characteristics, volatility information, and news and announcement semantics, within the same decision-making chain. By introducing a multi-agent collaborative mechanism, different functional agents can participate in trading decisions from the perspectives of trend perception, event judgment, risk identification, and execution control, thereby improving the completeness and stability of trading judgments under complex market conditions. Simultaneously, a large language model is used to extract key semantic clues from financial texts, enhancing the model's understanding of potential market disturbances and implicit trading signals. At the trading execution level, this paper further integrates explicit buy and sell conditions, as well as profit-taking and stop-loss constraints, to construct an adaptive trading mechanism for practical application scenarios. This ensures that the strategy output not only has strong information integration capabilities but also clear execution logic. The results show that the proposed method performs well in terms of profit generation, drawdown control, and overall transaction quality, verifying the effectiveness and application value of the constructed framework for automated stock trading in complex financial market environments.

Keywords: Automated stock trading; multi-source information fusion; collaborative decision-making mechanism; financial semantic modeling

1. Introduction

As the interconnectedness of global financial markets continues to increase, the external environment for stock trading activities is exhibiting complex characteristics of high-frequency fluctuations, multi-factor coupling, and rapid evolution [1, 2]. Macroeconomic changes, industry rotation, policy expectations, sudden event shocks, and the diffusion of investor sentiment are intertwined, resulting in significant uncertainty, non-stationarity, and dynamic correlation in market conditions. Against this backdrop, traditional stock trading methods based on single factors, static rules, or partial information modeling often fail to adequately characterize the heterogeneous information propagation mechanisms and time-varying decision-making logic in complex financial markets. Consequently, they are prone to under-adaptability, delayed response, and decreased robustness when facing drastic fluctuations, structural shifts, and information shocks. Constructing intelligent trading algorithms with dynamic perception, collaborative reasoning, and continuous adaptability for complex financial market environments has become a crucial issue in the field of financial intelligence research [3].

In recent years, the rapid development of artificial intelligence technology has provided new theoretical foundations and technical pathways for upgrading stock trading algorithms. Especially within a multi-agent collaborative decision-making framework, different functional agents can collaborate and divide labor around key aspects such as market perception, signal recognition, risk assessment, strategy generation, and execution control, thereby overcoming the limitations of single-agent models in terms of cognitive scope, decision-making depth, and environmental adaptability. Compared to traditional single-agent methods, multi-agent systems can better simulate the interactive decision-making process of multiple perspectives, levels, and objectives in complex markets, helping to improve the ability of trading strategies to handle market anomalies, information noise, and local distortions. Simultaneously, the information exchange and collaborative optimization mechanisms among multiple agents provide new research space for building a more flexible, robust, and globally consistent stock trading system [4].

Meanwhile, large language models have shown strong potential in semantic understanding, knowledge integration, contextual modeling, and complex task reasoning, bringing new enabling directions to the intelligentization of financial transactions. A large amount of key information in the stock market does not exist solely in structured numerical form but is widely distributed in policy texts, news information, company announcements, industry reports, public opinion content, and research viewpoints. This type of information possesses strong semantic meaning, timeliness, and implicit correlations, making it difficult for traditional quantitative models to achieve a deep understanding and effective utilization. Large language models can not only extract valuable semantic signals from massive amounts of unstructured financial text, but also provide stronger expressive capabilities and reasoning support in areas such as multi-source information fusion, event impact analysis, and strategy intent generation. Combining these models with multi-agent collaborative mechanisms holds promise for propelling stock trading algorithms from simply relying on numerical pattern mining to a stage of intelligent decision-making jointly driven by structured and semantic information.

Based on this, researching adaptive stock trading algorithms driven by multi-agent collaborative decision-making and large language models in complex financial market environments has significant theoretical and practical value. From a theoretical perspective, this research helps expand the construction paradigm of financial intelligent decision-making models, promotes the deep integration of multi-agent systems, large language models, and adaptive trading algorithms within a unified framework, and provides new research ideas for the analysis of intelligent decision-making mechanisms in complex dynamic environments [5]. From an application perspective, this research can support stock trading systems in achieving a higher level of environmental understanding, strategy optimization, and risk response under complex market conditions, helping to improve the sensitivity, adaptability, and sustainable decision-making capabilities of trading algorithms to market changes. Addressing the practical needs of the continuous evolution of financial technology and the intelligent transformation of capital markets, conducting research in this direction not only has a distinctly cutting-edge nature but also a strong practical driving force.

2. BackGround

As the digitalization of financial markets deepens, stock trading has evolved from traditional experience-based judgment and rule enforcement to a complex process relying on massive data analysis, dynamic information processing, and intelligent model decision-making. Modern stock markets are influenced not only by structured data such as prices, trading volumes, and technical indicators, but also by a combination of heterogeneous information from multiple sources, including macroeconomic policies, industry trends, corporate announcements, international events, and online public opinion. This market characteristic of diverse information sources, rapid dissemination, and complex impact paths makes stock price formation mechanisms more dynamic, interconnected, and uncertain [6]. Especially in complex financial market environments, market conditions can change significantly in a short period, and the effectiveness of trading signals can rapidly diminish with changing contexts, thus placing higher demands on the real-time perception, contextual understanding, and dynamic adjustment capabilities of trading algorithms. Traditional trading methods driven by fixed rules, single models, or local features are no longer sufficient to meet the real-world needs for efficient decision-making and flexible responses in complex market conditions.

Driven by the continuous development of artificial intelligence, research on stock trading algorithms is gradually shifting from single-model prediction to a comprehensive modeling approach that integrates multi-agent collaboration and semantically enhanced decision-making. On the one hand, multi-agent collaborative mechanisms can decompose tasks such as market analysis, risk identification, strategy selection, and trade execution, improving the overall adaptability and decision-making efficiency of the system through interaction and cooperation among different decision-making entities. On the other hand, large language models demonstrate significant advantages in complex text understanding, knowledge extraction, logical organization, and contextual reasoning, providing new technical conditions for the in-depth utilization of financial textual information [7]. In complex financial markets, many key influencing factors are often hidden in unstructured text, and relying solely on numerical sequences is insufficient to fully reveal their potential mechanisms. Therefore, combining the idea of multi-agent collaborative decision-making with the semantic modeling capabilities of large language models to construct an adaptive trading framework capable of simultaneously processing structured market signals and unstructured semantic information has become an important development direction in intelligent finance research and has also laid an important foundation for improving the decision-making quality and application value of stock trading algorithms in complex environments.

3. Methodology

To characterize the decision process of stock automatic trading under a complex financial environment, the proposed method formulates the market at time step as a hybrid observation system that jointly contains numerical signals, textual semantics, and agent interaction states. Let denote the structured market vector composed of price, volume, volatility, momentum, and liquidity indicators, while represents the semantic representation extracted from large language models over financial news, policy disclosures, company announcements, and sentiment streams. Beyond isolated feature extraction, the core motivation of this design is to prevent trading decisions from relying only on local price fluctuations, since short-term movements in complex markets are often entangled with latent macro signals and delayed semantic shocks. The overall architecture of the algorithm is further presented here, as shown in Figure 1.

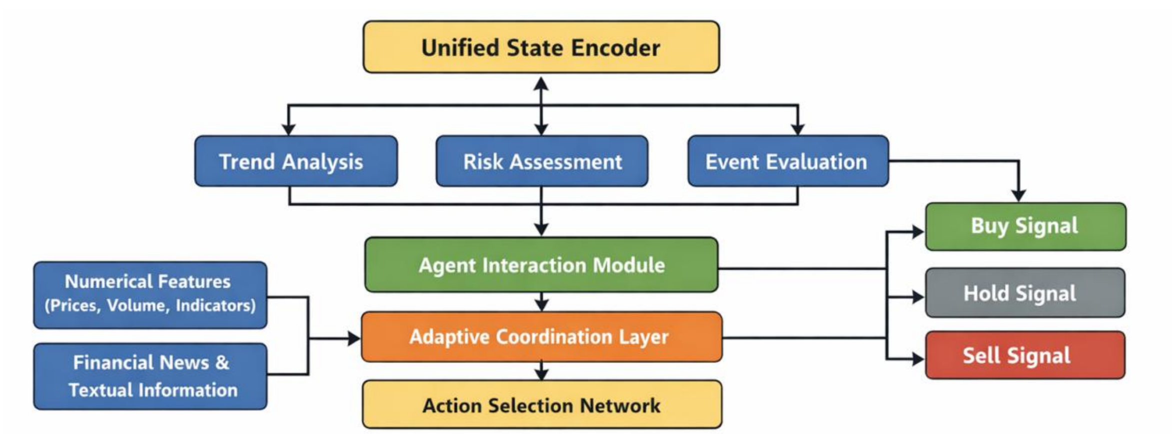

Figure 1. Overall model architecture

A unified state encoder is therefore introduced to compress heterogeneous observations into a compact latent representation that supports downstream coordination, risk control, and trade execution. The global market state is defined as:

$$\boldsymbol{s}_t = \phi(\boldsymbol{W}_m \boldsymbol{m}_t + \boldsymbol{W}_n \boldsymbol{n}_t + \boldsymbol{b}_s)$$

where denotes a nonlinear activation, and $\boldsymbol{W}_m$, $\boldsymbol{W}_n$, and are trainable parameters that align structured and unstructured information into the same decision space. Rather than assigning a single controller to process $\mathbf{s}_t$, the framework decomposes trading intelligence into multiple agents with different roles, including trend perception, event assessment, risk supervision, and execution timing, so that the resulting policy can better reflect the multi-perspective nature of real financial decision-making. Suppose that the hidden response of agent is written as:

$$\boldsymbol{h}_t^{(i)} = \sigma\left(\boldsymbol{W}_i \boldsymbol{s}_t + \sum_{j \neq i} \alpha_{ij,t} \boldsymbol{U}_{ij} \boldsymbol{h}_{t-1}^{(j)} + \boldsymbol{b}_i\right)$$

in which is the agent activation function, measures the influence from agent to agent at time t, and projects cross-agent messages into the latent space of the receiving agent. Such a coordination mechanism is meaningful because trading opportunities in turbulent markets are rarely determined by one source alone, and robust actions usually emerge only after market momentum, semantic context, and risk preference are jointly negotiated within the agent population.

Given the set of agent responses $\{\boldsymbol{h}_t^{(i)}\}_{i=1}^N$, an adaptive coordination layer aggregates them into a consensus representation that preserves both cooperation and specialization. This stage is introduced to avoid the common weakness of hard voting or naive averaging, because agents should not contribute equally when the market regime changes from stable drift to abrupt reversal. An attention-based fusion is thus defined as:

$$\beta_{i,t} = \frac{exp\left(\boldsymbol{q}^\top tanh\left(\boldsymbol{V}\boldsymbol{h}_t^{(i)}\right)\right)}{\sum_{k=1}^N exp\left(\boldsymbol{q}^\top tanh\left(\boldsymbol{V}\boldsymbol{h}_t^{(k)}\right)\right)}$$

and the coordinated decision state is computed by:

$$\boldsymbol{z}_t = \sum_{i=1}^N \beta_{i,t} \boldsymbol{h}_t^{(i)}$$

where and are trainable attention parameters, and dynamically reflects the relevance of each agent under the current market conditions. Meanwhile, a gating pathway is incorporated to regulate the influence of semantic information on the final trading representation, since textual evidence may be highly informative during event-driven periods yet potentially noisy during ordinary sessions. The gate value can be absorbed into the fusion process through trainable transformations on and $\boldsymbol{n}_t$, allowing the model to amplify macro-event awareness when semantic consistency is high and suppress language-derived disturbance when uncertainty rises. From a methodological perspective, this adaptive weighting strategy enables the trading system to remain flexible without sacrificing structural interpretability, which is crucial for financial decision chains requiring both responsiveness and internal coherence.

Under the coordinated latent state $\mathbf{z}_t$, the trading policy outputs continuous action scores for buying, holding, and selling, and these scores are further constrained by position status and risk exposure. A direct score projection is adopted not merely for classification convenience, but to preserve the smooth variation of action preference across adjacent market states and reduce unstable switching caused by marginal fluctuations. The raw policy output is formulated as:

$$\boldsymbol{a}_t = softmax(\boldsymbol{W}_a \boldsymbol{z}_t + \boldsymbol{b}_a)$$

with $\boldsymbol{a}_t = \left[a_t^{buy}, a_t^{hold}, a_t^{sell}\right]^{\top}$, where and map the latent coordination state into executable trading tendencies. Because automatic trading must explicitly define entry and exit conditions, the model sets the buy signal only when the current portfolio holds no position, the buy score dominates the other scores, and the expected upward return exceeds a risk-adjusted threshold. The entry rule is written as:

$$u_t^{buy} = \mathbb{I}(p_{t-1} = 0)\mathbb{I}\left(a_t^{buy} = max(\boldsymbol{a}_t)\right)\mathbb{I}(\hat{r}_{t+1} > \tau_b + \lambda_b \hat{\sigma}_t)$$

where denotes the previous position state, is the predicted next-period return, is the estimated short-horizon risk level, and together with controls the strictness of opening a long position. In contrast, the sell action is triggered when the portfolio is already invested and at least one exit principle becomes active, including expected downside deterioration, stop-loss breach, or take-profit achievement. The exit rule is defined as:

$$u_t^{sell} = \mathbb{I}(p_{t-1} = 1)\mathbb{I}\left(\hat{r}_{t+1} <- \tau_s \ \vee \ R_t^{unreal} \leq - \eta_l \ \vee \ R_t^{unreal} \geq \eta_p\right)$$

in which is the unrealized return of the open position, is the negative return tolerance, and and denote the stop-loss and take-profit boundaries. Through these definitions, trade execution becomes state-aware and rule-grounded at the same time, ensuring that buy and sell behaviors are not left vague but are instead embedded into the mathematical structure of the policy itself.

Finally, the portfolio state evolves according to the realized execution decision, making the proposed framework a closed-loop adaptive trading system rather than a disconnected predictor. Once the entry and exit indicators are determined, the position transition is updated by:

$$p_t = min\left(1, max\left(0, p_{t-1} + u_t^{buy} - u_t^{sell}\right)\right)$$

so that the holding status remains valid under mutually exclusive trading operations. Equally important, this transition equation ties decision logic to capital exposure, which gives the model a practical interpretation in stock automatic trading scenarios where whether to hold cash or maintain a position is itself part of intelligent control. Overall, the method integrates heterogeneous market sensing, multi-agent coordination, semantic reasoning from large language models, and explicit buy-sell execution rules into a unified adaptive framework. Such an organization is meaningful because it transforms stock trading from a narrow signal reaction problem into a structured collaborative decision problem, thereby strengthening the ability of the algorithm to respond to complex market variation with both informational breadth and policy discipline.

4. Experimental Results and Analysis

4.1 Dataset

This study constructs the dataset by taking all A-share stocks during the period from January 1, 2025, to December 31, 2025, as the initial research universe, so as to ensure that the samples cover trading behaviors, price fluctuation patterns, and information disturbance characteristics under different market stages throughout a complete year. In the process of data organization, the research scope includes the daily trading data and other available market information of ordinary listed stocks in the Shanghai and Shenzhen markets, and integrates price, trading activity, volatility, and auxiliary trading attributes into a unified data foundation. Such a data setting can more adequately reflect the dynamic characteristics of stock price evolution in complex financial market environments, while also providing continuous and representative sample support for subsequent multi-source information modeling, trading state characterization, and adaptive decision analysis.

To improve sample quality and enhance the consistency of the research objects, certain stocks were systematically excluded during the data screening process. First, small-cap stocks were removed because they usually exhibit insufficient liquidity, abnormally amplified price fluctuations, and stronger short-term trading noise, which can easily interfere with trading signal modeling and strategy stability. Second, stocks listed on the Beijing Stock Exchange were excluded because their trading mechanisms, liquidity structure, and market characteristics differ significantly from those of the main board and the ChiNext market, and retaining them may introduce distribution shifts caused by heterogeneous market rules. Finally, STAR Market stocks were further excluded because they possess strong particularities in listing standards, investor composition, price limit rules, and valuation characteristics, and their price formation mechanism is not fully consistent with that of general A-share stocks. After the above screening steps, a more stable, comparable, and automatic-trading-oriented A-share stock sample set was ultimately obtained.

4.2 Experimental Setup

This study configures and evaluates the proposed adaptive stock trading method under a unified experimental environment. The experimental subjects are a selected sample of A-share stocks for the entire year of 2025. The decision-making process is driven by both structured market data and textual semantic information, with GPT-5.4 serving as the core foundational model for semantic understanding and information enhancement modules in the construction of market semantic representation. At the transaction execution level, to ensure the consistency of strategy constraints and the clarity of risk control, both the profit-taking threshold and the stop-loss threshold are uniformly set at 3%, thus ensuring that the generation of buy and sell signals is not only affected by the comprehensive decision-making state but also constrained by the fixed return boundary. Meanwhile, the key configuration items involved in the experiment, such as the data time range, stock selection rules, core model, trading action constraints, and risk control parameters, are summarized in Table 1 to clearly illustrate the overall experimental setup.

Table 1. Experimental setup

Item	Setting
Market scope	A-share stocks
Time period	2025-01-01 to 2025-12-31
Initial universe	All A-share stocks within the study period
Excluded stocks	Small-cap stocks, Beijing Stock Exchange stocks, STAR Market stocks
Data type	Daily market data and available related market information
Core semantic model	GPT-5.4
Trading task	Adaptive stock automatic trading
Action space	Buy, Hold, Sell
Take-profit threshold	3%
Stop-loss threshold	3%
Trading objective	Joint optimization of adaptive decision making and risk-controlled execution

4.3 Experimental Results Compared with Other Models

To further illustrate the horizontal relationship between the proposed method and representative research in the same field, relevant literature on automated stock trading, multi-agent collaborative decision-making, reinforcement learning trading, and financial semantic modeling was selected as comparison objects, and a unified summary was made from the perspectives of profitability, risk-return balance, and drawdown control. Based on this, Table 2 summarizes the comparison of relevant methods on four indicators: ARR, SR, MDD, and CR.

Table 2. Comparison with representative related methods

Method	ARR	SR	MDD	CR
Lee et al. [8]	12.84	0.71	18.95	1.08
Yang et al. [9]	16.42	0.96	16.38	1.21
Liu et al. [10]	15.73	0.91	16.92	1.18
Zhang et al. [11]	18.65	1.08	15.44	1.29
Li et al. [12]	19.14	1.12	14.87	1.33
Fatemi et al. [13]	20.06	1.18	14.22	1.39
Zou et al. [14]	20.73	1.21	13.95	1.42
Ours	23.68	1.36	11.84	1.57

The algorithm proposed in this paper achieves a relatively harmonious balance between profit generation capability, trading stability, and risk control level, demonstrating strong comprehensive trading performance. This method can more effectively integrate multi-source market information and semantic information in complex financial market environments, improve the ability to identify key trading opportunities through a multi-agent collaborative decision-making mechanism, and enhance the robustness of strategy execution through adaptive decision-making and risk constraint mechanisms. Therefore, the constructed joint-driven framework not only helps improve the quality of returns in automated stock trading but also reduces drawdown pressure and improves overall capital utilization efficiency to a certain extent, indicating that the algorithm has good practical value and methodological effectiveness under complex market conditions.

4.4 Ablation Test Results

To further analyze the role of each key component in the overall automated trading framework, this paper conducts ablation analysis focusing on unified state coding, agent interaction mechanisms, and adaptive coordination mechanisms, and the relevant results are summarized in Table 3. By systematically eliminating core modules, the actual contributions of different functional units to profit generation, decision stability, and risk constraints can be more clearly observed, thereby verifying the structural rationality and synergistic effectiveness of the constructed method in complex financial market environments.

Table 3. Ablation experiment results

Ablation setting	ARR	SR	MDD	CR
w/o Unified State Encoder	17.94	1.03	15.62	1.28
w/o Agent Interaction Module	19.08	1.11	14.74	1.35
w/o Adaptive Coordination Layer	20.16	1.18	13.96	1.41
Ours	23.68	1.36	11.84	1.57

From an overall performance perspective, the algorithm proposed in this paper achieves a relatively coordinated and unified relationship between profit acquisition, risk constraints, and trading robustness, indicating that the constructed overall framework possesses strong structural integrity and decision-making effectiveness. This method does not rely on a single information source or local decision-making mechanism to complete trading decisions; instead, through the combined effects of unified state representation, multi-agent interaction, and adaptive coordination, it achieves a more comprehensive integration of key information in complex market environments. Therefore, the final model exhibits better overall performance in automated trading tasks, demonstrating that this joint modeling approach can effectively improve strategy quality and enhance the system's adaptability to complex market disturbances.

Furthermore, the key to the good results achieved by the proposed method lies in the functional interconnection and mutual support of its constituent modules. The unified state encoding mechanism strengthens the fusion capability of heterogeneous market information, providing a more stable input basis for the decision-making process; the multi-agent interaction mechanism enhances information collaboration between different decision-making perspectives, preventing trading decisions from being limited to a single path; and the adaptive coordination mechanism further improves the screening and integration capabilities of key decision signals, enabling the model to maintain good response quality when facing complex market changes. Therefore, the algorithm proposed in this paper has strong rationality and effectiveness in its overall design, and has good application value for adaptive stock trading tasks in complex financial market environments.

4.5 The Impact of Predictive Return Threshold Setting on MDD

The predicted return threshold directly affects the stringency of the buy trigger conditions, thus further altering the entry timing and subsequent drawdown control. Visualizing this parameter helps to more clearly characterize the risk exposure of the algorithm under different decision thresholds. Based on this setting, displaying the maximum drawdown performance under varying predicted return thresholds provides a more intuitive basis for parameter selection.

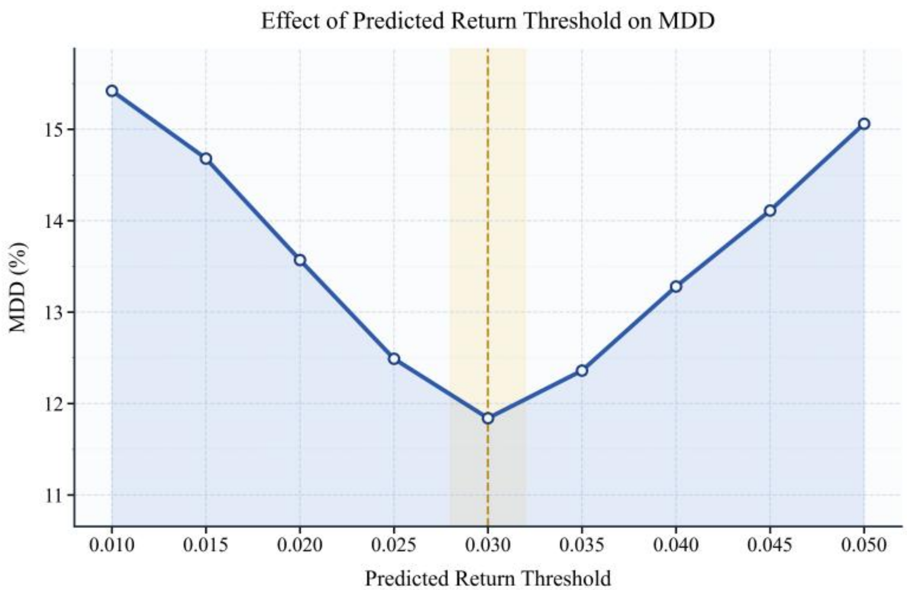

Figure 2. The impact of predictive return threshold setting on MDD

As shown in Figure 2, the algorithm proposed in this paper can effectively balance the effectiveness of trading decisions and the stability of risk control under these parameter settings, indicating that the design of the predicted return threshold plays a crucial role in the quality of strategy execution. A reasonable threshold constraint enables the model to more accurately screen effective buying opportunities in complex financial market environments, reducing ineffective trades triggered by low-quality signals, while avoiding additional drawdown pressure caused by overly lenient decision-making. This demonstrates that the proposed method possesses strong parameter adaptation and risk suppression capabilities in adaptive trading, further reflecting the practical value and effectiveness of the overall framework in automated stock trading tasks.

4.6 A Visualization Experiment on the Correlation between Trading Signal Confidence and Return Quality

The confidence level of trading signals reflects the strength of the algorithm's judgment on the current market state during automated trading decisions. Therefore, it is necessary to further examine its correlation with return quality. Visual analysis of this factor helps to understand the model's behavioral characteristics in complex financial market environments from the perspective of decision reliability. Based on this setting, the performance of return quality under changes in trading signal confidence level is displayed, providing a more intuitive reference for subsequent strategy explanation and parameter setting. The experimental results are shown in Figure 3.

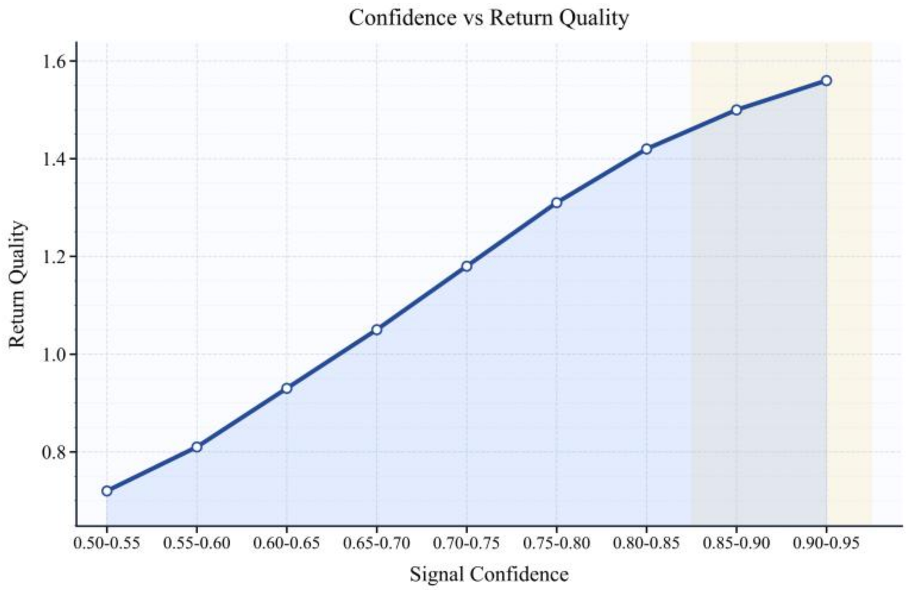

Figure 3. Visualization of the relationship between trading signal confidence and return quality under the proposed method.

Figure 3 illustrates that the proposed algorithm effectively transforms trading signal confidence into meaningful return quality, demonstrating strong consistency between decision-making and trade execution. The high-quality signal screening mechanism enables the strategy to more effectively identify valuable trading opportunities in complex financial market environments, while reducing the adverse impact of low-confidence decisions on overall return quality. This indicates that the proposed method not only possesses good decision reliability in automated trading tasks but also demonstrates strong return organization capabilities and practical application value.

5. Conclusion

This paper addresses the challenges of automated stock trading in complex financial market environments, including strong information heterogeneity, rapid market state changes, and limited single-path decision-making capabilities. It proposes an adaptive stock trading method driven by multi-agent collaborative decision-making and a large language model. Starting from the practical decision-making needs under complex market conditions, this research incorporates structured market data and unstructured semantic information into a unified modeling framework. Through unified state coding, multi-agent interaction and collaboration, and an adaptive coordination mechanism, a relatively complete intelligent trading chain is constructed, enabling the trading decision-making process to simultaneously possess market perception, semantic understanding, and risk constraint capabilities. Compared to traditional trading methods that rely on single numerical signals or static rules, this method places greater emphasis on the linkage between multiple sources of information in complex markets and the functional collaboration between decision-making entities, thus providing a more systematic, structured, and intelligent solution for automated stock trading tasks.

From an overall research value perspective, the significance of this work lies not only in the performance improvement at the trading algorithm level but also in the expansion of the intelligent decision-making paradigm in complex finance. By introducing large language models into automated stock trading scenarios, financial text, policy information, announcements, and market semantic signals can more effectively participate in the representation of trading states and decision generation. This allows automated trading systems to move beyond limited local numerical fluctuations such as price and transaction volume, enabling them to make strategy judgments within a broader information context. Simultaneously, the multi-agent collaborative mechanism further enhances the system's ability to organize market information from multiple perspectives, enabling closer linkages between trend identification, event assessment, risk control, and action selection. This methodological framework offers significant insights for promoting the development of applications in intelligent finance, quantitative investment, asset management, and algorithmic trading, and provides a valuable technical path for research on intelligent decision-making in complex dynamic environments. With the increasing demand for timely, robust, and adaptable decision-making systems in the financial market, the research approach proposed in this paper has good potential for widespread application.

Future research can be further deepened at several levels. Firstly, extended research can be conducted focusing on longer time spans, richer market structures, and more granular trading scenarios to further enhance the model's adaptability to cross-cycle market switching, sudden information shocks, and complex trading constraints. Secondly, based on the existing collaborative decision-making framework, the division of roles, interaction mechanisms, and dynamic game relationships among multiple agents can be further improved, enabling the system to more accurately depict the information dissemination and behavioral feedback processes involving multiple stakeholders in real financial markets. Thirdly, the semantic understanding depth and knowledge reasoning capabilities of large language models in professional financial contexts can be further enhanced, strengthening their integration efficiency with price and volume signals, time-series states, and risk constraints, thereby improving the stability, interpretability, and generalization ability of automated trading systems in complex environments. With the continuous development of artificial intelligence technology, financial information processing capabilities, and the demand for intelligent trading, research on adaptive stock trading driven by multi-agent collaborative decision-making and large language models is expected to unleash broader application value in areas such as intelligent investment advisory, institutional strategy support, market monitoring and early warning, and the construction of financial technology infrastructure platforms.

References

[1] Y. Xiao, E. Sun, D. Luo et al., "TradingAgents: Multi-Agents LLM Financial Trading Framework," arXiv preprint arXiv:2412.20138, 2024.

[2] J. A. Chudziak and M. Wawer, "ElliottAgents: A Natural Language-Driven Multi-Agent System for Stock Market Analysis and Prediction," Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation, pp. 961-970, 2024.

[3] Y. Lu, "Policy-Aware Enterprise Risk Assessment through LLM-Based Cross-Source Representation and Reasoning," 2024.

[4] C. Zhang, X. Liu, Z. Zhang et al., "When AI Meets Finance (StockAgent): Large Language Model-Based Stock Trading in Simulated Real-World Environments," arXiv preprint arXiv:2407.18957, 2024.

[5] S. Chen, H. Qiu, H. Huang and N. Eli, "Dynamic Low-Rank Routing for Efficient Multi-Task Instruction Fine-Tuning of Large Language Models," 2025.

[6] C. Zong, J. Wan, L. Cascone et al., "Stock Movement Prediction with Multimodal Stable Fusion via Gated Cross-Attention Mechanism," Complex & Intelligent Systems, vol. 11, no. 9, Art. no. 396, 2025.

[7] C. S. Lee, "Long-Range Dependency Modeling and Decision Point Summarization for Large Language Models in Dialogue and Meeting Scenarios," 2024.

[8] J. W. Lee and O. J., "A Multi-Agent Q-Learning Framework for Optimizing Stock Trading Systems," International Conference on Database and Expert Systems Applications, pp. 153-162, 2002.

[9] H. Yang, X. Y. Liu, S. Zhong et al., "Deep Reinforcement Learning for Automated Stock Trading: An Ensemble Strategy," Proceedings of the First ACM International Conference on AI in Finance, pp. 1-8, 2020.

[10] X. Y. Liu, H. Yang, Q. Chen et al., "FinRL: A Deep Reinforcement Learning Library for Automated Stock Trading in Quantitative Finance," arXiv preprint arXiv:2011.09607, 2020.

[11] H. Zhang, Z. Shi, Y. Hu et al., "Strategic Trading in Quantitative Markets through Multi-Agent Reinforcement Learning," arXiv preprint arXiv:2303.11959, 2023.

[12] Z. Li, J. Jiang, Y. Cao et al., "Logic-Guided Deep Reinforcement Learning for Stock Trading," The Second Tiny Papers Track at International Conference on Learning Representations, 2024.

[13] S. Fatemi and Y. Hu, "FinVision: A Multi-Agent Framework for Stock Market Prediction," Proceedings of the 5th ACM International Conference on AI in Finance, pp. 582-590, 2024.

[14] P. Zou and P. O. Siebers, "StockMARL: A Novel Multi-Agent Reinforcement Learning System to Dynamically Improve Trading Strategies," Procedia Computer Science, vol. 274, pp. 929-938, 2025.