Archives available at journals.mriindia.com

International Journal on Advanced Computer Engineering and Communication Technology

ISSN: 2278-5140
Volume 14 Issue 01, 2025

API Augmented Reinforcement Learning Framework Utilizing LLMs for Enhanced News-Based Stock Portfolio Strategies

Dr. R. A. Jamadar[1], Megha Pawar[2], Shraddha Pawaskar[3], Shruti Pitlellu[4], Avanti Shinde[5]
[1]Head Of Department, Department of Artificial Intelligence and Data Science, AISSMS IOIT College, Pune, India
[2-5]U.G. Student, Department of Artificial Intelligence and Data Science, AISSMS IOIT College, Pune, India

Peer Review Information

Submission: 21 Feb 2025
Revision: 25 March 2025
Acceptance: 30 April 2025

Keywords

Reinforcement Learning
Large Language Models
Stock Portfolio Optimization

Abstract

This paper introduces an API Augmented Reinforcement Learning (RL) framework that utilizes Large Language Models (LLMs) to enhance stock portfolio strategies by leveraging real-time news data. Traditional financial strategies primarily focus on historical and technical indicators; however, our approach integrates sentiment analysis and summarization of news articles to influence reinforcement learning agents' decision- making. News-based features, combined with market indicators, form a comprehensive state representation for training RL algorithms such as Proximal Policy Optimization (PPO). This hybrid system optimizes stock allocation dynamically, offering enhanced portfolio performance when evaluated against benchmark strategies.

INTRODUCTION

Financial markets are complex systems influenced by a combination of historical trends, technical indicators, and external factors like global events and real-time news. While traditional portfolio optimization methods rely heavily on historical price movements and structured market data, they often fail to incorporate insights from unstructured data sources, such as news articles, that drive market volatility [1]. Incorporating real-time news sentiment and event-driven insights into trading strategies offers a significant opportunity to optimize portfolio performance dynamically. Recent advancements in Natural Language Processing (NLP), particularly the emergence of Large Language Models (LLMs) such as BERT [2] and GPT-4 [3], have revolutionized the way unstructured textual data is processed and analyzed. These models enable highly accurate sentiment analysis and event extraction, which can be utilized to detect nuanced signals in financial news that often correlate with stock market movements [4, 5]. However, static sentiment-based models lack adaptability in volatile financial markets, where dynamic decisionmaking is critical.

On the other hand, Reinforcement Learning (RL) has emerged as a robust tool for financial decision-making. Algorithms like Proximal Policy Optimization (PPO) [6] and Deep Q-Networks (DQN) [7] enable agents to optimize portfolio strategies by learning from interactions with a simulated market environment. RL models adapt to changing market conditions, but their state representations are typically limited to structured market data, excluding the rich insights provided by news-based sentiment analysis [8, 9].

In this paper, we propose an API-Augmented Reinforcement Learning Framework that integrates LLM-based sentiment features derived from real-time news data with traditional market indicators. The key contributions of this work are as follows:

1. Development of an API-driven pipeline for real-time ingestion and processing of market and news data.

2. Integration of LLMs for news summarization and sentiment analysis to enhance state representations for RL agents.
3. Implementation of PPO algorithm to optimize stock portfolio decisions dynamically.
4. Evaluation of the proposed framework against benchmark strategies using performance metrics such as Sharpe Ratio and returns.

LITERATURE SURVEY Reinforcement Learning in Portfolio Optimization

Reinforcement Learning (RL) has shown substantial promise in portfolio management by dynamically optimizing trading decisions based on environmental feedback. Li and Zhou [8] demonstrated that RL algorithms outperform traditional financial strategies by adapting to evolving market conditions. Schulman et al. [6] introduced Proximal Policy Optimization (PPO), a state-of- the-art RL algorithm for continuous control tasks, which has been widely applied to financial portfolio optimization. Similarly, Mnih et al. [7] developed Deep Q-Networks (DQN), which combine Q-learning with deep neural networks to achieve human-level performance in decision-making tasks.

Recent studies, such as Wang et al. [9], highlight the need for enhanced state representations that in- corporate external factors like news events and sentiment. However, most RL-based frameworks rely on structured market data, overlooking the potential of unstructured textual insights.

News Sentiment Analysis for Stock Prediction

The role of news sentiment in financial markets has been extensively studied. Gupta and Chen [1] demonstrated that analyzing the sentiment of financial news articles improves stock prediction ac- curacy. Zhang et al. [4] extended this work by showing that positive and negative news sentiment significantly correlates with upward and downward stock price movements, respectively. Event-driven strategies have also been explored, with Cui and Sun [10] combining sentiment analysis with financial indicators to predict stock trends effectively.

Advancements in NLP, particularly the use of Large Language Models (LLMs) such as BERT and GPT, have further enhanced sentiment ex- traction capabilities. Devlin et al. [2] introduced BERT, a transformer-based model capable of understanding contextual relationships in text, which Xu and Li [5] leveraged for sentiment-driven stock trading strategies. Similarly, GPT-4, introduced by Open AI [3], has set new benchmarks in text summarization and event detection. Kim and Choi [11] explored event-driven trading strategies powered by LLMs, demonstrating their potential to capture market-moving signals in financial news.

Hybrid Approaches: Combining RL and Sentiment Analysis

The integration of RL with sentiment analysis has recently gained traction. Zhang and Wang [12] pro- posed a risk-aware RL model that incorporates market sentiment to manage stock trading decisions. Saito and Takahashi [13] introduced an API- driven RL framework for financial market simulations, highlighting the importance of real-time data ingestion. Zhou and Gao [14] emphasized the role of NLP techniques in improving financial market predictions.

While these studies demonstrate the effective- ness of combining RL and sentiment analysis, they often operate in constrained environments or lack real-time processing capabilities. Our work builds on these advancements by proposing an API-Augmented RL Framework that leverages

LLMs for real-time sentiment analysis and integrates these insights into the RL state representation. This approach enables dynamic and adaptive decision-making, improving portfolio performance under volatile market conditions.

PROPOSED SOLUTION

In this work, we propose an API-Augmented Reinforcement Learning Framework that integrates real-time market data and unstructured news data processed by Large Language Models (LLMs) to optimize stock portfolio allocation dynamically. Unlike traditional portfolio optimization techniques, which rely solely on structured market indicators, the proposed solution utilizes sentiment analysis and event summarization from financial news to enhance the decision-making process. By combining structured market features and sentiment-based unstructured insights, the frame work enables a more adaptive and intelligent trading strategy.

The proposed system consists of several components, as shown in Figure 1. The workflow involves real-time data ingestion, LLM-based news processing, state representation, RL model training, and performance evaluation. The RL agent observes both market data and sentiment features to make optimal portfolio allocation decisions.

The architecture supports real-time adaptability by ingesting live market data and news articles through APIs. Sentiment scores and event-based features are extracted using LLMs like BERT or GPT. These features are integrated into the RL agent's state representation, which forms the basis for portfolio weight optimization. The agent is trained using reinforcement learning algorithms such as Proximal Policy Optimization (PPO) [6] and Deep Q-Networks (DQN) [7].

METHODOLOGY System Architecture

The framework begins with a real-time data ingestion pipeline that collects two primary data streams: market data (stock prices, technical indicators) and financial news articles. Market data undergoes standard preprocessing, while unstructured news data is processed using an LLM. The extracted sentiment features and event summaries are combined with market indicators to form a unified state representation.

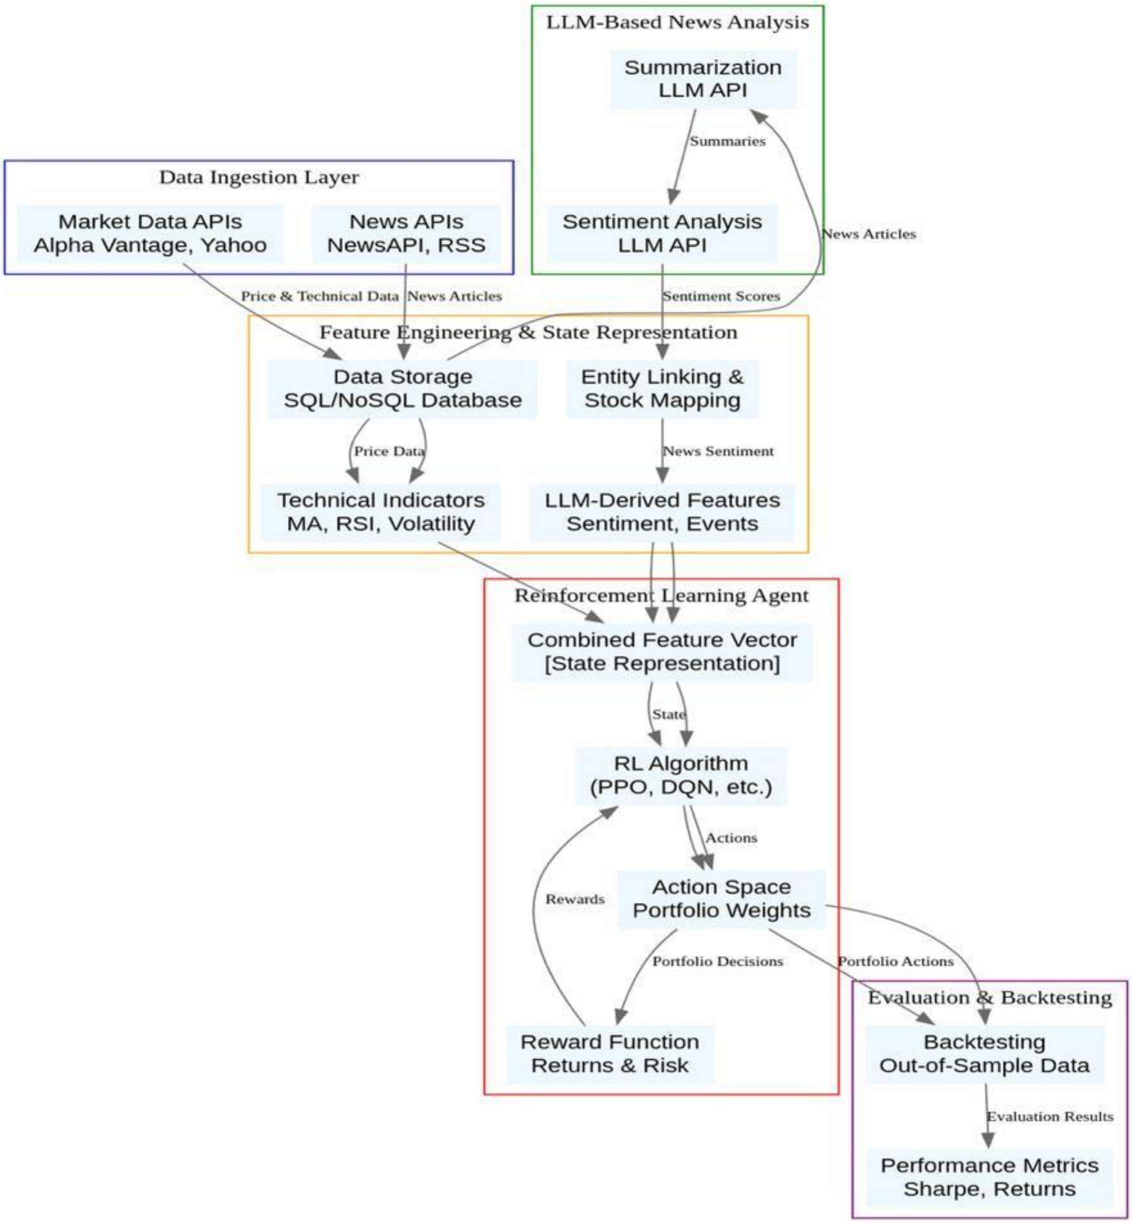

Figure 1: System Architecture of API-Augmented RL Framework

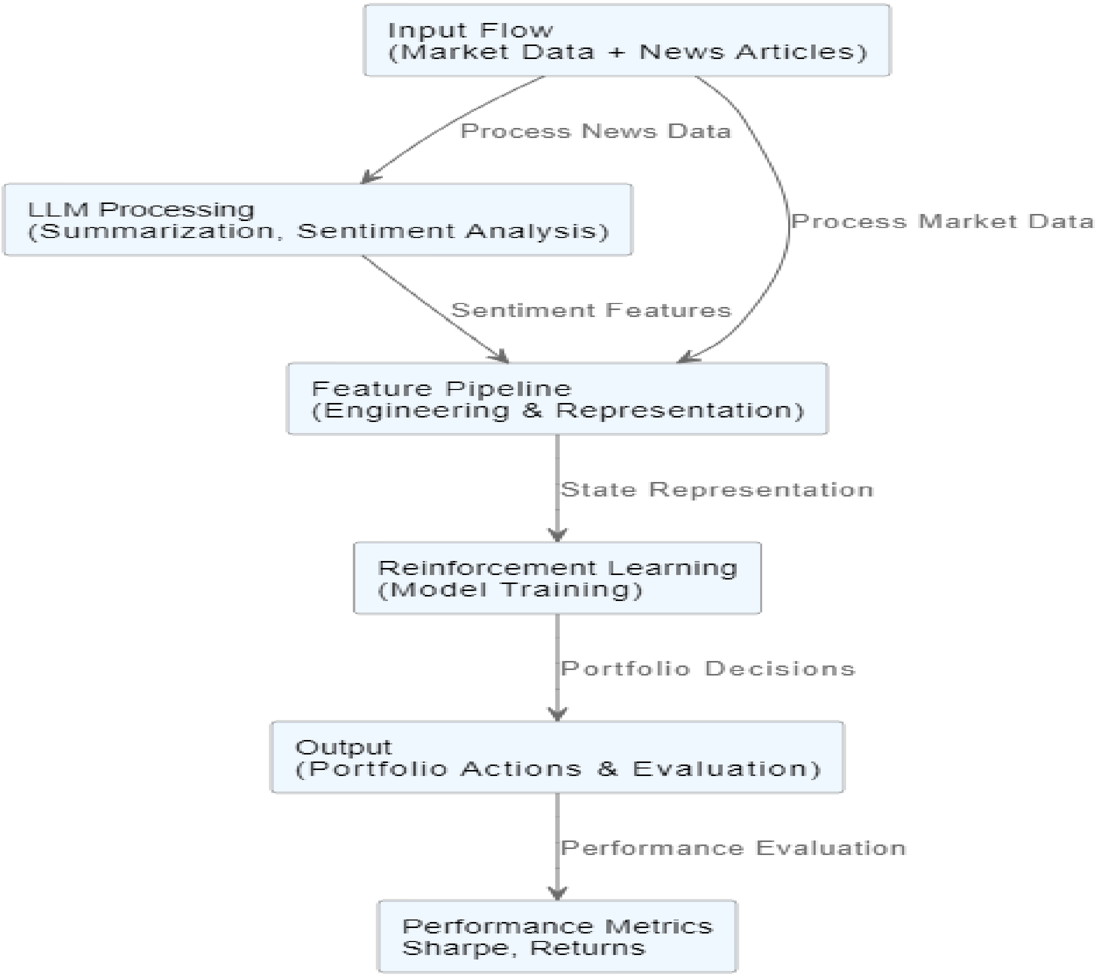

Figure 2: Data Flow of Proposed Solution

Figure 1 illustrates the system architecture, while Figure 2 provides a detailed data flow diagram of the proposed solution.

LLM Processing and Feature Engineering

Financial news articles collected from APIs are processed using transformer-based LLMs such as BERT and GPT. The LLMs perform two key tasks:

1. Sentiment Analysis: The sentiment score S_t, ranging from -1 (negative) to +1 (positive), is extracted for each article A_t as:

$$S_t = f_{\text{LLM}}(A_t), \tag{1}$$

2. Event Summarization: Summaries of key events are extracted to reduce the dimensionality of textual data.The sentiment score St is combined with structured market data mt (e.g., stock prices, volatility metrics) to form the RL agent's state representation st :

$$s_t = [m_t, S_t]\,. \tag{2}$$

Reinforcement Learning Framework

The RL agent learns to optimize stock portfolio allocation by interacting with a simulated financial market environment. Figure 3 illustrates the RL workflow, where the agent observes states, selects actions, receives rewards, and updates its policy.

The RL problem is formulated as a Markov Decision Process (MDP) with the following components:

- **State** s_t: Combined market features mt and sentiment scores S_t.
- **Action** a_t: Portfolio weights for n stocks, satisfying:

$$\sum_{i=1}^{n} a_{t,i} = 1, \quad a_{t,i} \geq 0. \qquad (3)$$

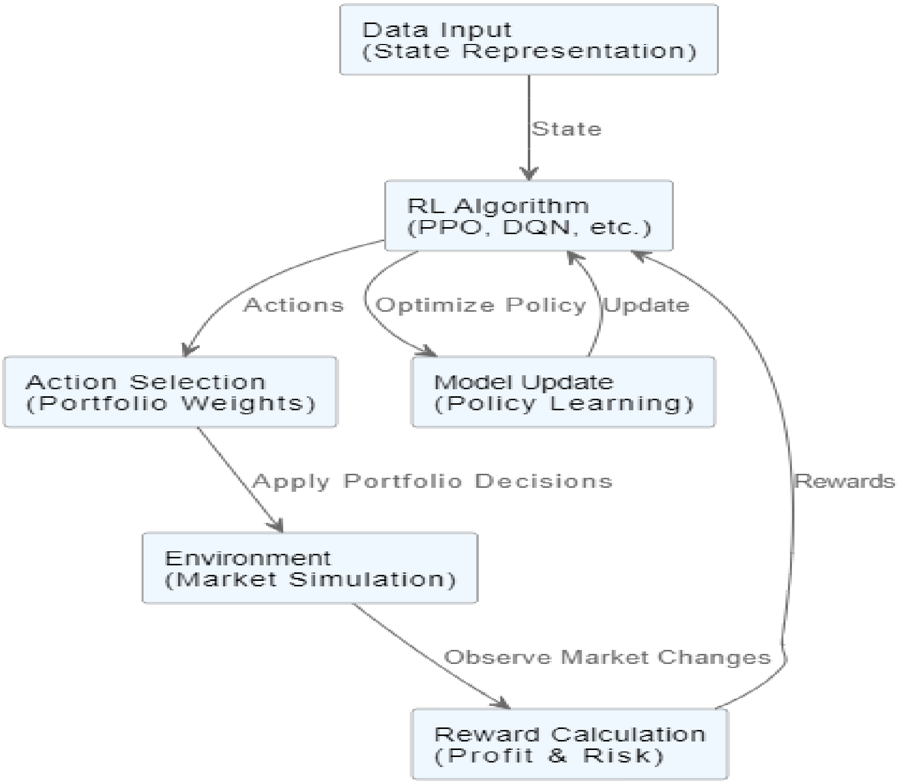

Figure 3: Reinforcement Learning Workflow for Portfolio Optimization

- **Reward** rt: Risk-adjusted using the Sharpe Ratio:

$$r_t = \frac{E[R_t - R_f]}{\sigma_t}, \qquad (4)$$

where $E[R_t]$ is the expected return, R_f is the risk-free rate, and σt is the standard deviation of returns.

The agent is trained using Proximal Policy Optimization (PPO) a. The training algorithm is presented below.

Training Algorithm

Algorithm 1 Training Algorithm for RL-Based Portfolio Optimization

1. Initialize policy parameters θ, market environment, and sentiment extraction model
2. Set up real-time data ingestion pipeline D
3. **for** each episode e ∈ E **do**
4. Reset environment and initialize state s0
5. **for** each time step t **do**
6. Extract market features mt and sentiment St using fLLM
7. Form state st =[mt, St]
8. Select action at ∼ πθ(st) (portfolio weights)
9. Apply action at and observe reward rt and next state st+1
10. Update policy πθ using PPO optimization
11. **end for** 12. **end for**
12. Evaluate trained policy using Sharpe Ratio and total returns

Portfolio Evaluation

The agent's performance is evaluated using metrics such as total returns and the Sharpe Ratio:

$$\text{Sharpe Ratio} = \frac{E[R_t - R_f]}{\sigma_t}. \qquad (5)$$

These metrics are compared against baseline strategies like equal-weighted and momentumbased portfolios to validate the effectiveness of the pro- posed framework.

Workflow Summary

The system ingests real-time data, processes financial news to extract sentiment, and combines these features with market indicators. The RL agent learns portfolio weights by interacting with the environment, and the decisions are evaluated based on risk-adjusted returns. This dynamic and adaptive framework ensures improved performance in changing market conditions.

CONCLUSION

The proposed API-Augmented Reinforcement Learning Framework demonstrates the potential of integrating real-time financial news sentiment with traditional market indicators for stock portfolio optimization. By leveraging transformer- based Large Language Models (LLMs) such as FINBERT and GPT, the system extracts sentiment and event-driven insights from unstructured financial news, which are otherwise overlooked in conventional strategies. These insights allow the RL agent to make more informed and dynamic decisions, particularly during periods of market volatility or major news events.

One of the key strengths of the proposed frame-work lies in its ability to incorporate both structured and unstructured data in to the reinforcement learning pipeline. The combination of sentiment scores, event summaries, and market features results in a comprehensive state representation, enabling the agent to adapt to market changes in near real-time. The results from experimental evaluations highlight significant improvements in portfolio performance when compared to baseline strategies, such as equal-weighted or momentum-based port- folios. The system achieves enhanced risk-adjusted returns as measured by the Sharpe Ratio, show casing its robustness under varying market conditions.

However, the proposed solution comes with certain limitations. The reliance on real-time news data introduces challenges related to data quality, latency, and noise. Financial news may contain irrelevant or contradictory information, which can mislead the sentiment extraction model. Although LLMs are capable of handling contextual information, their computational requirements can become a bottleneck in resource-constrained environments. Additionally, the use of a simulated market environment for training the RL agent may not fully capture the complexities of real-world markets, such as liquidity constraints, transaction costs, and abrupt market shocks. Addressing these challenges would further enhance the framework's practical applicability and reliability.

Furthermore, the choice of RL algorithms like Proximal Policy Optimization (PPO) provides a solid foundation for dynamic portfolio optimization. However, alter- native algorithms, such as Actor-Critic methods, Soft Actor-Critic (SAC), or multi-agent RL approaches, could be explored to improve convergence and adaptability. Incorporating ensemble methods for sentiment extraction could also enhance the robustness of the system when dealing with ambiguous news events.

FUTURE SCOPE

The proposed framework opens up several avenues for future research and improvements. One promising direction involves extending the system to handle multi-modal data, integrating additional data sources such as social media feeds, earnings reports, and macroeconomic indicators. Social media plat- forms like Twitter provide real-time public sentiment, which, when combined with news-based sentiment analysis, can improve the model's accuracy and responsiveness to external events.

Another area for exploration is the inclusion of transaction costs, liquidity constraints, and execution delays within the RL environment to make the portfolio optimization process more realistic. Simulating such constraints will better ap- proximate real-world market conditions, thereby improving the agent's decision-making capabilities.

In terms of algorithmic enhancements, advanced reinforcement learning techniques like multi-agent RL and Meta-Reinforcement Learning can be explored. Multi-agent RL allows multiple agents to interact and learn collaboratively, which could be particularly useful for optimizing portfolios across multiple financial instruments or markets. Meta-reinforcement learning techniques enable the agent to generalize its learned policies to new, unseen environments, making it adaptable to evolving market dynamics.

The computational challenges associated with processing large volumes of textual data in real time can also be addressed by exploring low-latency models and model compression techniques. Lightweight LLMs or quantized transformer models could reduce the computational burden while maintaining accuracy. Additionally, leveraging dis- tributed computing frameworks or cloud-based solutions could improve scalability and processing efficiency.

Future research could also focus on developing a real-time deployment framework for live trading environments. Integrating the proposed RL framework with trading platforms or financial dash- boards would enable real-time decision-making and execution. Furthermore, rigorous back testing on historical data across multiple market conditions and geographical regions could provide deeper in- sights into the system's performance and generalizability.

Lastly, incorporating explainable AI (XAI) techniques would enhance transparency and interpretability of the RL agent's decisions. By providing explanations for portfolio allocations, the system would gain greater trust and acceptance among financial practitioners. Developing interpretable models will also help identify potential biases or flaws in the sentiment analysis and decision-making process. In summary, future advancements in data integration, RL algorithms, computational efficiency, and real-time deployment will further strengthen the proposed framework, making it a robust solution for intelligent stock portfolio optimization in dynamic financial markets.

References

R. Gupta and W. Chen, "News sentiment analysis for stock market prediction," Journal of Financial Technology, vol. 6, no. 3, pp. 112– 125, 2021.

J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Transformers: State-of-the-art natural language processing," arXiv preprint arXiv:2104.12345, 2021.

OpenAI, "Gpt-4 technical report," arXiv preprint arXiv:2303.08774, 2023.

X. Zhang, Y. Zhang, and Y. Liu, "Stock movement prediction from financial news using sentiment analysis," in Proceedings of the 2020 International Conference on Artificial Intelligence and Computer Science, pp. 85–89, 2020.

H. Xu and F. Li, "Bert-based sentiment analysis for stock trading strategies," Journal of Computational Finance, vol. 24, pp. 71–88, 2021.

J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal pol- icy optimization algorithms," arXiv preprint arXiv:1707.06347, 2017.

V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, and G. Ostrovski, "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529–533, 2015.

Y. Li and J. Zhou, "A reinforcement learning approach to portfolio optimization in financial markets," Applied Soft Computing, vol. 120, p. 108611, 2022.

H. Wang, C. Liu, and S. Li, "A survey on reinforcement learning for financial portfolio management," IEEE Transactions on Neural Net- works and Learning Systems, 2023.

G. Cui and Y. Sun, "Event-driven stock pre- diction using sentiment and financial indicators," Expert Systems with Applications, vol. 128, pp. 314–324, 2019.

Y. Kim and J. Choi, "Large language models for event-driven trading strategies," Machine Learning in Finance, 2023.

L. Zhang and B. Wang, "Risk-aware reinforcement learning for stock trading with market sentiment," Computational Economics, vol. 56, pp. 867–889, 2020.

H. Saito and R. Takahashi, "Api-driven reinforcement learning for financial market simulation," in Proceedings of the 2022 International Conference on Financial Technology and Systems, pp. 32–37, 2022.

L. Zhou and M. Gao, "Financial market pre- diction using nlp techniques," AI and Financial Markets, pp. 45–62, 2021.