

Dynamic Strategy Optimization for Algorithmic Trading in the Indian Stock Market Using Reinforcement Learning and Sentiment Analysis

Vikram Singh Rathore¹, Dr. Riya Mukhopadhyay²,

¹*Masters of Technology, PES University, Hosur Rd, Konappana Agrahara, Bangalore*

²*Assistant Professor, CMS JAIN(Deemed-to-be-University)*

Abstract- Technical aspects and current market mood have a significant impact on the extremely volatile Indian stock market. Conventional static trading techniques, like Moving Average Crossover, RSI, and MACD-based methods, frequently perform below par since they are unable to adjust to abrupt changes in the market. In order to improve trading performance and decision-making, this study suggests a novel dynamic strategy optimisation framework that combines sentiment analysis and reinforcement learning (RL). Our approach combines sentiment scores obtained from financial news and social media utilising FinBERT with historical market data, such as price and volume indices. Based on sentiment patterns and current market conditions, the machine learning model is trained to assess and choose the trading strategy that performs the best. The system optimises risk-adjusted returns by constantly adjusting trading strategies through the combination of technical and sentiment analysis. Backtesting using actual Indian stock market data is used to evaluate performance, looking at important parameters including maximum drawdown, Sharpe ratio, and total return. The findings show that trading performance is much improved by sentiment integration, particularly during erratic market events such as earnings reporting or budget announcements. This study advances financial technology and intelligent trading systems by offering a scalable, automated method for strategy optimisation.

key words: Algorithmic Trading, Reinforcement Learning, Sentiment Analysis, Technical Indicators, Indian Stock Market, Machine Learning, FinBERT, Dynamic Strategy Optimization.

INTRODUCTION

Technical indicators and current market sentiment both contribute to the dynamic nature of the Indian stock market. Conventional algorithmic trading systems frequently perform poorly in turbulent environments since they mainly rely on static

technical indicators like Moving Averages, RSI (Relative Strength Index), and MACD (Moving Average Convergence Divergence). These methods result in less-than-ideal decisions and more risk since they are unable to adjust to the quick changes brought about by sentiment shifts from social media and financial news. This study offers a fresh solution to these problems by combining sentiment analysis and reinforcement learning (RL) to dynamically optimise trading tactics. In order to optimise performance and control risk, the system adjusts to shifting market conditions by fusing technical indicators with real-time sentiment scores generated by models such as FinBERT. In order to automate strategy selection and offer better informed, data-driven buy/sell recommendations, this method makes use of sophisticated machine learning algorithms. To train the RL model, the process entails gathering and examining historical stock market data in addition to sentiment trends. Based on real-time inputs, the model dynamically chooses the best trading strategy after learning to assess a variety of methods, such as moving average crossovers, RSI-based, and MACD-based strategies. A strong, flexible system that enhances risk-adjusted returns in the erratic Indian stock market is the end product. By providing a novel, scalable approach to dynamic strategy optimisation, this work bridges the gap between technical and sentiment-based trading systems and advances the field of financial technology.

LITERATURE REVIEW

By combining generative models, deep reinforcement learning (DRL), and machine learning (ML), algorithmic trading has advanced significantly. Though recent advances in computational finance provide intriguing alternatives, traditional trading strategies frequently

find it difficult to adjust to dynamic and volatile markets. Deep Reinforcement Learning (DRL) has become a leading approach for optimizing trading strategies. Zhang et al. [67] investigated DRL techniques, including Deep Q-Networks (DQN), Policy Gradients (PG), and Advantage Actor-Critic (A2C), to improve trading performance for futures contracts. Utilizing discrete and continuous action spaces alongside volatility scaling, their results demonstrated that DRL-based methods outperformed traditional time-series momentum models, even under high transaction cost scenarios. Similarly, Thate et al. [53] introduced the Trading Deep Q-Network (TDQN), designed to maximize Sharpe ratios using prioritized experience replay and artificial trajectory generation. Their findings showed significant advantages over conventional buy-and-hold strategies, especially in volatile markets such as BTC/USD. Generative Adversarial Networks (GANs) have also been employed to address challenges in algorithmic trading. Sun et al. [49] proposed a GAN-based solution to mitigate backtest overfitting by generating synthetic market scenarios that mimic real-world conditions. This approach enhanced the robustness and out-of-sample performance of trading strategies. Similarly, Wang et al. [55] developed Factor-GAN, which integrates GANs with multi-factor asset pricing models. The proposed model exhibited exceptional accuracy and stability in predicting stock returns under volatile conditions, achieving higher Sharpe ratios compared to traditional models. Transformer-based architectures have shown considerable potential in financial modeling. Mai et al. [37] introduced StockGPT, a generative pretrained transformer designed specifically for stock prediction. The model achieved impressive results, with an annualized return of 119 DRL advancements have also extended to portfolio management. Rheya et al. [45] proposed an ensemble strategy combining multiple DRL algorithms, such as Proximal Policy Optimization (PPO) and Soft ActorCritic (SAC), to adapt to varying market conditions. This method delivered improved risk-adjusted returns and lower volatility compared to traditional approaches. Building on this, Yu et al. [64] presented dynamic decision strategies, including Nested Reinforcement Learning and Weighted Random Selection with Confidence, to optimize agent selection and trading outcomes in diverse market environments. The combination of cascaded LSTM networks with DRL has proven effective for managing noisy and volatile

financial data. Zou et al. [68] introduced an LSTM-PPO framework for automated stock trading, achieving notable improvements in cumulative returns and risk management across emerging and established markets. Broader reviews of algorithmic trading and machine learning applications have further highlighted their impact on market behavior. Damilare et al. [13] discussed the dual role of algorithmic trading in enhancing market efficiency while contributing to volatility, underscoring the need for adaptive regulatory frameworks. Similarly, Addy et al. [1] conducted a critical review of machine learning applications in trading and risk management, emphasizing their advantages in predictive modeling and signal generation while addressing challenges such as overfitting and ethical considerations. 2 These studies collectively demonstrate the progress in algorithmic trading, highlighting the effectiveness of DRL, GANs, and transformers in adapting to complex market environments. Future research avenues include integrating macroeconomic factors, overcoming computational challenges, and improving model interpretability to further advance these technologies.

MATERIALS AND METHODS

Data Collection

Source The stock and announcement data used in this project were sourced from the National Stock Exchange (NSE) of India (<https://www.nseindia.com/>). The datasets included: Historical equities data: Important trading details including volumes and prices. Announcements from companies: events that are pertinent to the market, like board meetings, earnings reports, and company activities.

Data Points

Fields like Open, High, Low, Close, Volume, and other trading-related variables that are essential for examining market behaviour were included in the equity data. Announcements: Contains information that may affect market sentiment, such as the date, time, and content of the announcement. Sentiment Scores: Generated from announcement data using FinBERT, a financial sentiment analysis model, and classifying each announcement as either positive, neutral, or negative.

Download Process

To guarantee the data's legitimacy and dependability, it was taken straight from the NSE website. To confirm the integrity and completeness of the data, additional manual checks were carried out.

Data Preprocessing

Date and Format Standardization Dates across datasets were standardized to the YYYY-MM-DD format to ensure seamless merging and alignment. This step was critical to maintaining temporal consistency in the analysis.

Column Cleanup **Cleaning Column Names:** Extra spaces and inconsistencies in column names were removed. **Dropping Irrelevant Columns:** Fields such as ATTACHMENT, DIFFERENCE, DISSEMINATION, and RECEIPT, which were not required for analysis, were removed to streamline the dataset.

Sentiment Analysis for Announcements

Announcements in the DETAILS column were categorised as positive, neutral, or negative using FinBERT, a cutting-edge model trained for financial mood. A sentiment score that measured the polarity of the sentiment was given to each announcement, offering a useful tool for trading strategy optimisation and further study.

Handling Missing Data Numerical Data Cleaning:

To ensure correct numerical processing, commas were eliminated from numerical columns labelled Open, High, Low, Close, and Volume. **Imputation of Missing Values:** To ensure data correctness and continuity, missing values were handled by linear interpolation.

Final Processed Data

The final dataset consisted of pre-processed equity and announcement data, enriched with sentiment scores. It was merged across all selected stocks into a unified dataset, with a SYMBOL column added to identify each stock. The processed data was saved as a CSV file (final merged equity announcements with sentiment.csv) for subsequent model training and analysis.

Key Insights

- **Market Trends:** Patterns in price movements were captured using technical indicators such as Moving Averages, RSI, and MACD.

- **Sentiment Trends:** Positive and negative sentiment scores correlated with stock price volatility and direction, especially during critical events like earnings reports and budget announcements.

The integration of sentiment scores and technical indicators enabled the creation of a robust dataset capable of supporting machine learning and reinforcement learning models for dynamic trading strategy optimization.

INTERMEDIATE RESULTS AND DISCUSSION

Datasets In this project, the datasets are pivotal for training and evaluating the trading strategies. The data includes historical stock price information, financial news sentiment, and volume metrics. Key aspects of the datasets are:

- **Equity Data:** Historical price data (Open, High, Low, Close, Volume) for selected stocks from sources like NSE.

- **Sentiment Data:** Sentiment scores for financial news, which are used to gauge market sentiment and enhance trading decisions.

- **Processed Dataset:** A merged dataset including both price data and sentiment indicators, normalized and pre-processed for optimal model performance.

Performance Metrics

The performance of the trading models is assessed through several metrics, providing insights into their accuracy, efficiency, and overall effectiveness.

Accuracy of Detection

This metric measures the accuracy of the model's ability to detect favorable trading signals based on historical data and sentiment analysis. Detection accuracy is calculated by comparing predicted signals with actual price movements.

Accuracy of Identification

Identification accuracy refers to the model's precision in categorizing various market conditions, such as trend_following or mean reversion. This accuracy helps in assessing how well the model adapts to different market scenarios.

Time

Time metrics track the computation duration for processing data and executing trades. These include:

- **Training Time:** The time required for training this 10 to 12 hour.

- Inference Time: The time taken by the model to make predictions or execute trades in real-time in very fast took a few minutes

RESULTS

The following results demonstrate the performance of various trading strategies, including the Adaptive Strategy Model, based on key metrics such as Sharpe Ratio, Max Drawdown, and Cumulative Return.

Performance Metrics of Trading Strategies Table 1 provides a summary of the key performance metrics for all tested strategies, highlighting their risk adjusted returns, maximum drawdowns, and cumulative returns.

Strategy	Sharpe Ratio	Max Drawdown (%)	Cumulative Return (%)
Mean Reversion	0.57	-99.18	617.15
Volatility Breakout	0.56	-97.31	630.77
Momentum	-0.06	-97.36	-85.94
Buy and Hold	0.41	-99.88	-98.11
Trend Following	0.56	-100	-99.81
Adaptive Strategy Model	-0.55	-99.41	-98.32

Table 1: Comparison of Performance Metrics for Trading Strategies.

Net Worth Over Time

Figures 1 to 6 illustrate the net worth trends over time for each strategy, highlighting their performance dynamics:

- Mean Reversion and Volatility Breakout: Both strategies achieved significant cumulative gains, demonstrating strong adaptability to market fluctuations.
- Momentum and Buy and Hold: These static strategies suffered heavy losses due to their inability to adjust to volatile conditions.
- Adaptive Strategy Model: Despite its adaptability framework, the strategy underperformed, with a sharp decline in net worth over time (Figure 6).

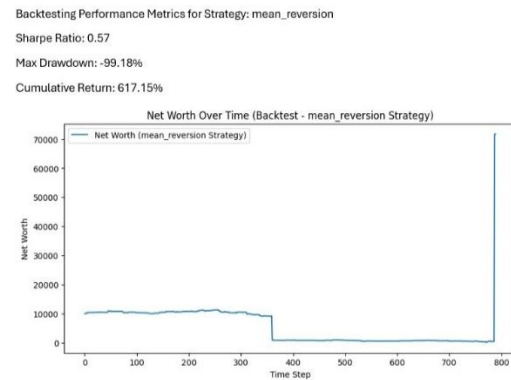


Figure 1: Net Worth Over Time for Mean Reversion Strategy

DISCUSSION OF RESULTS

Key insights from the performance analysis include:

- Top Performers: The Mean Reversion and Volatility Breakout strategies achieved cumulative returns of 617.15% and 630.77%, respectively, with Sharpe Ratios above 0.5, highlighting their effectiveness.
- Underperforming Strategies: The Adaptive Strategy Model, with a Sharpe Ratio of -0.55 and a cumulative return of -98.32%, failed to demonstrate profitability, despite its adaptability.
- Risk Management: Strategies with robust risk controls, such as Mean Reversion and Volatility Breakout, outperformed those lacking such mechanisms (e.g., Trend Following, Adaptive Strategy).
- Time Efficiency: The Adaptive Strategy Model exhibited reasonable training and inference times but suffered from performance instability, indicating the need for further optimization.

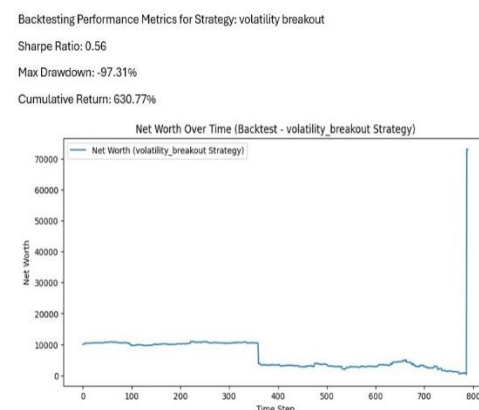


Figure 2: Net Worth Over Time for Volatility Breakout Strategy

Backtesting Performance Metrics for Strategy: momentum

Sharpe Ratio: -0.06

Max Drawdown: -97.36%

Cumulative Return: -85.94%

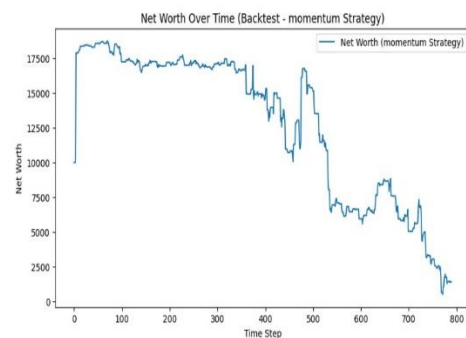


Figure 3: Net Worth Over Time for Momentum Strategy

Recommendations for the Adaptive Strategy Model

While the Adaptive Strategy Model shows potential, significant refinements are needed:

- **Hyperparameter Tuning:** Adjusting learning rate, entropy coefficients, and actor-critic losses could improve stability.
- **Risk Controls:** Incorporating stop-loss mechanisms and optimizing trade entry/exit criteria may reduce drawdowns and enhance returns.
- **Hybrid Models:** Combining adaptive strategies with proven approaches like Volatility Breakout could improve overall robustness.

Backtesting Performance Metrics for Strategy: buy_and_hold

Sharpe Ratio: 0.41

Max Drawdown: -99.88%

Cumulative Return: -98.11%

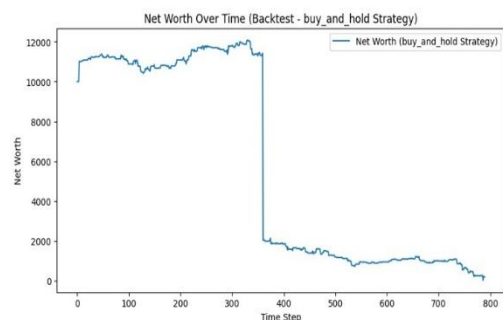


Figure 4: Net Worth Over Time for Buy and Hold Strategy.

Backtesting Performance Metrics for Strategy: trend_following

Sharpe Ratio: 0.56

Max Drawdown: -100.00%

Cumulative Return: -99.81%




Figure 5: Net Worth Over Time for Trend Following Strategy.

CONCLUSION

The Backtesting results provide valuable insights into the effectiveness of various trading strategies under different market conditions. The Mean Reversion and Volatility Breakout strategies emerged as the top performers, achieving cumulative returns of 617.15% and 630.77%, respectively, alongside strong Sharpe Ratios (0.57 and 0.56). These results highlight their ability to adapt to market dynamics and capture profitable opportunities.

On the other hand, the Buy and Hold and Momentum strategies showed significant underperformance, with cumulative returns of -98.11% and -85.94% and poor risk-adjusted performance. The Buy and Hold strategy suffered particularly from prolonged drawdowns, underscoring its vulnerability in volatile markets. Similarly, the Trend Following and Risk-Adjusted strategies faced extreme losses, with drawdowns nearing 100%, emphasizing the importance of robust risk management.

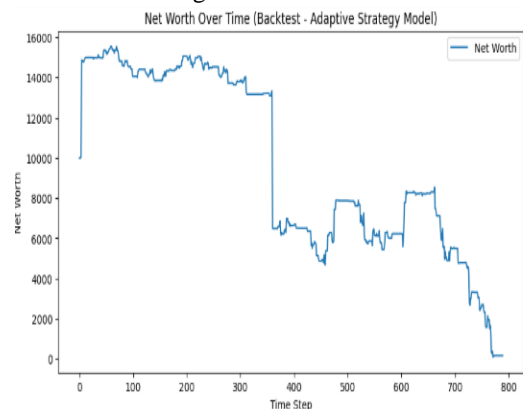


Figure 6: Net Worth Over Time for Adaptive Strategy Model

Risk management mechanisms, such as stop-loss and take-profit levels, played a pivotal role in mitigating large losses and stabilizing performance. Strategies lacking these measures suffered from heightened volatility and poor returns. Future enhancements could involve developing hybrid models that combine the strengths of multiple strategies, incorporating sentiment analysis and external market factors to improve predictive capabilities, and refining risk management systems to minimize drawdowns while maximizing returns.

REFERENCES

- [1] Wilhelmina Afua Addy, Adeola Olusola Ajayi-Nifise, Binaebi Gloria Bello, Sunday Tubokirifuruar Tula, Olubusola Odeyemi, and Titilola Falaiye. Machine learning in financial markets: A critical review of algorithmic trading and risk management. *International Journal of Science and Research Archive*, 2024.
- [2] Et. al. Ai Rosita. Market price signal prediction based on deep learning algorithm. *Turkish Journal of Computer and Mathematics Education (TURCOMAT)*, 2021.
- [3] Anonymous. A comprehensive survey of financial machine learning techniques. *International Journal of Financial Management and Research (IJFMR)*, 11(1):1–15, 2024.
- [4] Anonymous. Financial news dataset for stock price movement prediction. *arXiv preprint arXiv:2403.02647*, 2024.
- [5] Dogu Araci. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*, 2019.
- [6] Samuel A. Assefa. Generating synthetic data in finance: opportunities, challenges and pitfalls. *Proceedings of the First ACM International Conference on AI in Finance*, 2020.
- [7] Swagatam Chakraborty, Saroj Kumar Das, Helena Krammer, and Tirthankar Chakraborty. Financial time series forecasting using deep learning: A survey. *Soft Computing*, 26(3):1–29, 2022.
- [8] S. Kumar Chandar. Convolutional neural network for stock trading using technical indicators. *Automated Software Engineering*, 29, 2022.
- [9] Chang Che, Zengyi Huang, Chen Li, Haotian Zheng, and Xinyu Tian. Integrating generative ai into financial market prediction for improved decision making, 2024.
- [10] Lin Chen and Qiang Gao. Application of deep reinforcement learning on automated stock trading. 2019 *IEEE 10th International Conference on Software Engineering and Service Science (ICSESS)*, pages 29–33, 2019.
- [11] Ayantse Cornelius, Yaya S. OlaOluwa, Ojo F. Johnson, Akinlana M. Damola, and Oguntola Toyin Omoyeni. Trading volume, intraday and overnight volatility of stock returns in the nigerian stock exchange. *Asian Journal of Probability and Statistics*, 2024.
- [12] C.Saravanan. Application of artificial intelligence (ai) in forecasting share prices in stock markets a conceptual study. *International Journal of Creative Research Thoughts (IJCRT)*, ISSN:2320-2882, Volume.9, Issue 9, pp.115-122, September 2021, 2021.
- [13] Lawrence Damilare, Lawrence Damilare Oyeniyi, Chinonye Esther Ugochukwu, and Noluthando Zamanjomane Mhlongo. Analyzing the impact of algorithmic trading on stock market behavior: A comprehensive review. *World Journal of Advanced Engineering Technology and Sciences*, 2024.
- [14] Wenqi Fan, Shijie Wang, Jiani Huang, Zhikai Chen, Yu Song, Wenzhuo Tang, Haitao Mao, Hui Liu, Xiaorui Liu, Dawei Yin, et al. Graph machine learning in the era of large language models (llms). *arXiv preprint arXiv:2404.14928*, 2024.
- [15] Georgios Fatouros, Konstantinos Metaxas, John Soldatos, and Dimosthenis Kyriazis. Can large language models beat wall street? unveiling the potential of ai in stock selection. *arXiv preprint arXiv:2401.03737*, 2024.
- [16] Tianhao Gao, Yuwen Huo, and Mengmeng Yu. The progress of stock price forecasting model based on ai techniques: Arima, neural networks and lstm. *EAI*, 2 2023.
- [17] James Garza. Reinforcement learning for stock option trading. In 2023 31st Irish Conference on Artificial Intelligence and Cognitive Science (AICS), pages 1–4. *IEEE*, 2023.
- [18] Roberto Gozalo-Brizuela and Eduardo C. Garrido-Merch'an. A survey of generative ai applications, 2023.
- [19] Ganesh Harke, Mikhail Shishlenin, and Suresh Koppiseti. Application of algorithmic trading strategies for retail investors. Available at SSRN 3904097, 2020.
- [20] Xuan He and Ying He. Market requirements document on the use of ai to facilitate and

- manage stock price prediction. Highlights in Science, Engineering and Technology, 85:904–908, 2024. 9
- [21] Badr Hirchoua, Brahim Ouhbi, and Bouchra Frikh. Deep reinforcement learning based trading agents: Risk curiosity driven learning for financial rules-based policy. *Expert Systems with Applications*, 170:114553, 2021.
- [22] Gang Hu. Advancing algorithmic trading: A multi-technique enhancement of deep q-network models, 2023.
- [23] Chien Yi Huang. Financial trading as a game: A deep reinforcement learning approach. *arXiv preprint arXiv:1807.02787*, 2018.
- [24] Rahul Jain and Rakesh Vanzara. Emerging trends in ai-based stock market prediction: A comprehensive and systematic review. *Engineering Proceedings*, 56(1):254, 2023.
- [25] Tanveer Kabir, Ismail Ahmed, and Tahsin Sultana. Survey paper on stock prediction using machine learning algorithms. *International Journal of Creative Research Thoughts (IJCRT)*, 9(5):3311–3316, 2021.
- [26] Sana Kadri and Souheyl Rahali. Financial market prediction using machine learning algorithms: A review. *Journal of Artificial Intelligence and Intelligent Systems*, 2(1):29–41, 2021.
- [27] Minseok Kong and Jungmin So. Empirical analysis of automated stock trading using deep reinforcement learning. *Applied Sciences*, 13(1), 2023.
- [28] Adriano S. Koshiyama, Nick Firoozye, and Philip C. Treleaven. Generative adversarial networks for financial trading strategies fine-tuning and combination. *CoRR*, abs/1901.01751, 2019.
- [29] Hemanth P Kumar and Basavaraj S Patil. Forecasting volatility trend of inr usd currency pair with deep learning lstm techniques. In *2018 3rd International Conference on Computational Systems and Information Technology for Sustainable Solutions (CSITSS)*, pages 91–97. IEEE, 2018.
- [30] Ming Che Lee, Jia-Wei Chang, Sheng-Cheng Yeh, Tsorng-Lin Chia, Jie-Shan Liao, and Xu-Ming Chen. Applying attention-based bilstm and technical indicators in the design and performance analysis of stock trading strategies. *Neural Computing & Applications*, 34:13267 – 13279, 2022.
- [31] Hui Li, Yunpeng Cui, Shuo Wang, Juan Liu, Jinyuan Qin, and Yilin Yang. Multivariate financial time-series prediction with certified robustness. *IEEE Access*, 8:109133–109143, 2020.
- [32] Zechu Li, Xiao-Yang Liu, Jiahao Zheng, Zhaoran Wang, Anwar Walid, and Jian Guo. Finrl-podracar: high performance and scalable deep reinforcement learning for quantitative finance. In *Proceedings of the second ACM international conference on AI in finance*, pages 1–9, 2021.
- [33] Xiao-Yang Liu, Ziyi Xia, Jingyang Rui, Jiechao Gao, Hongyang Yang, Ming Zhu, Christina Wang, Zhaoran Wang, and Jian Guo. Finrl-meta: Market environments and benchmarks for data-driven financial reinforcement learning. *Advances in Neural Information Processing Systems*, 35:1835–1849, 2022.
- [34] Xiao-Yang Liu, Hongyang Yang, Qian Chen, Runjia Zhang, Liuqing Yang, Bowen Xiao, and Christina Dan Wang. Finrl: A deep reinforcement learning library for automated stock trading in quantitative finance. *arXiv preprint arXiv:2011.09607*, 2020.
- [35] Zimo Liu, Lin Yang, and Tami Takada. Predicting stock prices using tweet frequency and ai: Leveraging social media insights to forecast tomorrow’s market trends. In *CS & IT Conference Proceedings*, volume 13. CS & IT Conference Proceedings, 2023.
- [36] Lennart Ljung. Financial time series forecasting using support vector machines. *Physica A: Statistical Mechanics and its Applications*, 338(1-2):363–372, 2004.
- [37] Dat Mai. Stockgpt: A genai model for stock prediction and trading. *ArXiv*, abs/2404.05101, 2024.
- [38] N Maknickiene, I Lapinskaite, and A Maknickas. Application of ensemble of recurrent neural networks for forecasting of stock market sentiments. *equilib qj econ policy*. 2018; 13: 7–27, 2018.
- [39] Lingyun Mao and Jianping Qian. An overview of machine learning techniques in financial fraud detection. *Journal of King Saud University-Computer and Information Sciences*, 34(11):1356–1372, 2022.
- [40] Tamerlan Mashadihasanli. Stock market price forecasting using the arima model: an application to istanbul, turkiye. *Journal of*

- Economic Policy Researches, 9(2):439–454, 2022.
- [41] Giuseppe Masi, Matteo Prata, Michele Conti, Novella Bartolini, and Svitlana Vyetrenko. On correlated stock market time series generation. In Proceedings of the Fourth ACM International Conference on AI in Finance, pages 524–532, 2023. 10
- [42] Wassim El-Hajj Mariam Mokalled and Mohamad Jaber. Automated stock price prediction using machine learning. In Proceedings of the Second Financial Narrative Processing Workshop (FNP 2019), pages 16–24, 2019.
- [43] Joseph Tafataona Mtetwa, Kingsley Ogudo, and Sameerchand Pudaruth. Explainable algorithmic trading: Unlocking the black box with gan-based visualizations. In 2024 International Conference on Artificial Intelligence, Big Data, Computing and Data Communication Systems (icABCD), pages 1–10, 2024.
- [44] Joseph Tafataona Mtetwa, Kingsley A. Ogudo, and Sameerchand Pudaruth. Explainable algorithmic trading: Unlocking the black box with gan-based visualizations. 2024 International Conference on Artificial Intelligence, Big Data, Computing and Data Communication Systems (icABCD), pages 1–10, 2024.
- [45] Bhooshi Uhitha Rheya. Reinforcement learning for stock portfolio management. International Journal of Creative Research Thoughts (IJCRT), 10(8):d342–d351, August 2022.
- [46] Dev Shah, Haruna Isah, and Farhana Zulkernine. Stock market analysis: A review and taxonomy of prediction techniques. International Journal of Financial Studies, 7(2):26, 2019.
- [47] Jingyi Shen and M. Omair Shafiq. Short-term stock market price trend prediction using a comprehensive deep learning system. Journal of Big Data, 2020.
- [48] Yikun Shen, Xiaodan Jia, Tianyu Feng, and Nenghai Yu. Machine learning in finance: Overview, applications and challenges. Advances in Computational Economics, 40:229–256, 2019.
- [49] Ao Sun and Yuh-Dauh Lyuu. Backtesting trading strategies with gan to avoid overfitting. ArXiv, abs/2209.04895, 2022.
- [50] Steven Sun, Lingyun Mao, Yifan Wei, Ke Hu, Peiran Liu, Tianwei Zhang, and Wei Li. Finbert: A pretrained language model for financial communications. arXiv preprint arXiv:1908.10063, 2019.
- [51] Md Arif Istiaque Sunny, Mirza Mohd Shahriar Maswood, and Abdullah G Alharbi. Deep learning-based stock price prediction using lstm and bi-directional lstm model. In 2020 2nd novel intelligent and leading emerging sciences conference (NILES), pages 87–92. IEEE, 2020.
- [52] Richard S Sutton and Andrew G Barto. Reinforcement learning: An introduction. Robotica, 17(2):229–235, 1999.
- [53] Thibaut Th'eate and Damien Ernst. An application of deep reinforcement learning to algorithmic trading. Expert Systems with Applications, 173:114632, July 2021.
- [54] Xingchen Wan, Jie Yang, Slavi Marinov, Jan-Peter Calliess, Stefan Zohren, and Xiaowen Dong. Sentiment correlation in financial news networks and associated market movements. Scientific reports, 11(1):3062, 2021.
- [55] Jiawei Wang and Zhen Chen. Factor-gan: Enhancing stock price prediction and factor investment with generative adversarial networks. PLOS ONE, 19, 2024.
- [56] Yongfeng Wang and Guofeng Yan. Survey on the application of deep learning in algorithmic trading. Data Science in Finance and Economics, 1(4):345–361, 2021.
- [57] Magnus Wiese, Lianjun Bai, Ben Wood, and Hans Buehler. Deep hedging: Learning to simulate equity option markets, 2019.
- [58] GWRI Wijesinghe and RMKT Rathnayaka. Stock market price forecasting using arima vs ann; a case study from cse. In 2020 2nd International Conference on Advancements in Computing (ICAC), volume 1, pages 269–274. IEEE, 2020.
- [59] Phil Winder. Reinforcement learning. O'Reilly Media, 2020.
- [60] Jheng-Long Wu, Xian-Rong Tang, and Chin-Hsiung Hsu. A prediction model of stock market trading actions using generative adversarial network and piecewise linear representation approaches. Soft Computing, 27(12):8209–8222, 2023.
- [61] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. Bloomberggpt:

- A large language model for finance. arXiv preprint arXiv:2303.17564, 2023. 11
- [62] Tiancheng Xia and Zhidao Chen. Stock price anomalies analysis and detection based on machine learning. Highlights in Business, Economics and Management, 2023.
- [63] Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. Fingpt: Open-source financial large language models. arxiv. arXiv preprint arXiv:2306.06031, 2023.
- [64] X. Yu, W. Wu, and X. Liao. Dynamic stock-decision ensemble strategy based on deep reinforcement learning. Applied Intelligence, 53:2452–2470, January 2023. Accepted: 09 April 2022, Published: 09 May 2022, Issue Date: January 2023.
- [65] Frensi Zejnullahu, Maurice Moser, and Joerg Osterrieder. Applications of reinforcement learning in finance– trading with a double deep q-network. arxiv. arXiv preprint arXiv:2206.14267, 2022.
- [66] Jin Zhang and Dietmar G. Maringer. Indicator selection for daily equity trading with recurrent reinforcement learning. Proceedings of the 15th annual conference companion on Genetic and evolutionary computation, 2013.
- [67] Zihao Zhang, Stefan Zohren, and Stephen J. Roberts. Deep reinforcement learning for trading. In The Journal of Financial Data Science, 2019.
- [68] Jie Zou, Jiashu Lou, Baohua Wang, and Sixue Liu. A novel deep reinforcement learning based automated stock trading system using cascaded lstm networks. ArXiv, abs/2212.02721, 2022.