Asia Pacific Science Press

Asia Pacific Economic and Management Review
ISSN(O):3005-9275
ISSN(P):3005-9267
Vol. 2 No.2 (2025)

Deep Learning for Stock Performance Prediction: A Sharpe Ratio-Optimized Approach

Markus Schäfer[1], Thomas Müller[1], Dr. Hugo Fernandez[2]*

1.RWTH Aachen University, Germany

2.University of Barcelona, Spain

Corresponding author: Dr. Hugo Fernandez

Abstract: Accurate stock performance prediction is critical for portfolio management, risk assessment, and algorithmic trading. Traditional forecasting models often focus on minimizing prediction error but fail to consider risk-adjusted returns, making them suboptimal for real-world investment applications. Recent advances in deep learning have significantly improved financial time series forecasting, yet existing models primarily optimize for accuracy rather than maximizing risk-adjusted performance metrics such as the Sharpe ratio.

This study proposes a Sharpe ratio-optimized deep learning framework for stock performance prediction, integrating risk-sensitive forecasting mechanisms directly into model training. By embedding Sharpe ratio-based loss functions, the model prioritizes investment strategies that yield higher returns per unit of risk. The framework utilizes temporal convolutional networks (TCNs) and attention-based transformers, allowing for both short-term price trend detection and long-range dependency modeling. Additionally, reinforcement learning is employed to dynamically adjust portfolio allocation strategies based on evolving market conditions, ensuring adaptability across different asset classes.

Empirical results on real-world stock market datasets demonstrate that the proposed model outperforms traditional forecasting approaches in both predictive accuracy and financial performance. The study highlights the importance of integrating risk-sensitive optimization techniques within deep learning-based stock prediction frameworks, offering a more practical and scalable solution for quantitative investment strategies.

Keywords: Stock Performance Prediction; Deep Learning; Sharpe Ratio Optimization; Risk-Aware Forecasting; Portfolio Management; Reinforcement Learning

Published: Mar 20, 2025

DOI: https://doi.org/10.62177/apemr.v2i2.210

1. Introduction

Stock market prediction is a fundamental aspect of financial analysis, influencing investment strategies, risk management, and asset allocation decisions. Accurate forecasting enables traders and portfolio managers to anticipate market movements, optimize trade execution, and mitigate financial risks. However, stock price movements are inherently volatile, influenced by complex interactions between macroeconomic factors, investor sentiment, and liquidity conditions[1-5]. Traditional statistical models such as autoregressive integrated moving average (ARIMA) and generalized autoregressive conditional heteroskedasticity (GARCH) have been widely used in time series forecasting but are limited by their assumptions of linearity and stationarity [6]. These methods often struggle to adapt to dynamic market conditions and fail to capture nonlinear dependencies in financial data.

Deep learning has emerged as a powerful alternative, offering models that can learn hierarchical patterns from large-scale financial datasets [7]. Recurrent neural networks (RNNs) and long short-term memory (LSTM) networks have demonstrated success in capturing sequential dependencies within stock price data, improving predictive accuracy over traditional methods. However, RNN-based models are constrained by their sequential processing nature, which makes them computationally expensive and limits their ability to handle long-range dependencies effectively[8]. Transformer-based architectures, which leverage self-attention mechanisms, have addressed these challenges by allowing models to process entire time series in parallel while preserving temporal dependencies. These models have been widely adopted for financial forecasting, achieving state-of-the-art performance in stock trend prediction and volatility modeling.

Despite these advancements, most deep learning-based stock prediction models focus solely on minimizing forecast error using loss functions such as mean squared error (MSE) or mean absolute error (MAE) [9]. While this approach improves predictive accuracy, it fails to account for the financial implications of investment decisions. In real-world applications, investors prioritize risk-adjusted returns rather than raw prediction accuracy[3]. Traditional forecasting methods do not incorporate financial performance metrics such as the Sharpe ratio (SR), which measures the return per unit of risk. As a result, stock predictions generated by these models may not align with investment objectives, leading to suboptimal portfolio allocations [10].

SR is a widely used performance metric in portfolio management, assessing the trade-off between risk and return. Conventional forecasting models evaluate SR as a post-processing step rather than integrating it directly into model training [11-14]. This study proposes a deep learning framework that optimizes for SR during training, ensuring that the model's forecasts contribute to enhanced financial performance. The proposed approach incorporates temporal convolutional networks (TCNs) and transformer-based architectures to capture both short-term price fluctuations and long-term market trends. By embedding SR-based constraints within the model's loss function, the framework prioritizes predictions that lead to improved risk-adjusted returns[4].

In addition to SR optimization, the framework integrates reinforcement learning (RL) techniques to dynamically adjust forecasting strategies based on market conditions. Traditional deep learning models rely on static training data and require frequent retraining to adapt to new market regimes. RL enables the model to learn optimal decision-making policies, adjusting its forecasting thresholds in response to evolving risk-reward dynamics. This adaptability ensures that the model remains effective across different market environments, from stable trends to high-volatility periods.

The proposed framework is evaluated using historical stock market datasets, including price movements, trading volume, and macroeconomic indicators. The model's performance is compared against baseline approaches such as LSTMs and standard transformer-based predictors. Experimental results demonstrate that the SR-optimized deep learning model outperforms conventional forecasting techniques in both predictive accuracy and financial performance. The findings highlight the importance of integrating risk-sensitive optimization within deep learning-based stock prediction models, offering a practical and scalable solution for quantitative trading and portfolio management.

2. Literature Review

Stock performance prediction has been extensively studied in both academic research and financial industry applications. Traditional forecasting models have primarily relied on statistical time series methods, while the emergence of deep learning-based approaches has significantly improved predictive performance. Despite these advancements, existing models largely focus on minimizing prediction error, neglecting the integration of risk-adjusted performance metrics, which are critical for real-world financial decision-making[15]. This section reviews conventional statistical forecasting methods, machine learning-based models, transformer-based approaches, and the role of risk-aware optimization in financial forecasting.

Early forecasting techniques were built on statistical models such as ARIMA and its variants, which assume linear dependencies between historical and future values [16-20]. These models effectively capture stationary trends but struggle with nonlinearity and sudden shifts in financial markets [21]. GARCH models extended these capabilities by incorporating time-varying volatility, making them useful for risk estimation [22]. However, these methods rely on strong statistical assumptions that limit their adaptability in highly dynamic and complex financial environments [23]. Markets frequently experience abrupt changes due to macroeconomic factors, investor sentiment shifts, and global events, making it difficult for these models to generalize across different market regimes.

Machine learning techniques introduced nonparametric models capable of capturing complex relationships within financial data [24-27]. Approaches such as support vector machines and random forests improved predictive accuracy by learning nonlinear patterns [28]. However, these models do not inherently account for sequential dependencies in financial time series, as they treat observations as independent data points rather than as part of a continuous sequence. To address this limitation, RNNs and LSTMs were introduced, offering improved sequential modeling through memory-based learning. LSTMs demonstrated superior performance over traditional models by retaining information over extended time horizons, capturing both short-term fluctuations and long-term market trends. Despite their advantages, LSTMs and other RNN-based architectures suffer from vanishing gradient problems, limiting their effectiveness when processing long sequences[29].

Transformer-based models have emerged as a superior alternative, overcoming the scalability limitations of recurrent architectures [30]. Unlike RNNs, transformers process entire time series in parallel using self-attention mechanisms, making them particularly effective for long-range dependency modeling. This capability allows transformers to dynamically assign importance to different time steps, improving the model's ability to detect significant price movements. Studies have shown that transformers outperform both traditional deep learning architectures and conventional forecasting methods in financial applications, achieving superior accuracy in stock price prediction, volatility forecasting, and market trend analysis[31]. However, despite their advancements in predictive accuracy, transformer models still focus primarily on minimizing MSE or MAE, rather than optimizing for financial objectives such as SR.

Risk-adjusted metrics such as SR are essential for evaluating the trade-off between return and volatility [32-35]. Despite their significance in portfolio management, most deep learning-based forecasting models do not integrate SR into their optimization processes. Instead, SR is typically calculated as a post-processing evaluation metric rather than being embedded within the training objective. This approach results in models that generate accurate forecasts but do not necessarily align with investment strategies that prioritize risk-adjusted returns [36-40]. Recent research has explored ways to incorporate financial risk metrics into deep learning architectures, demonstrating that embedding risk-sensitive constraints can improve both predictive robustness and real-world financial applicability. However, most of these approaches remain limited to external risk constraints rather than fully integrating SR optimization into model training [41-44].

The proposed framework addresses this limitation by embedding SR optimization directly within the deep learning architecture. Unlike conventional forecasting techniques that focus solely on minimizing error metrics, this approach ensures that stock performance predictions align with investment objectives by explicitly optimizing for risk-adjusted returns. By incorporating SR constraints into the loss function, the model learns to prioritize forecasts that maximize return efficiency while minimizing downside risk[9]. Reinforcement learning techniques further enhance adaptability, allowing the model to dynamically adjust risk preferences and forecasting thresholds based on evolving market conditions.

Integrating risk-sensitive optimization into transformer-based stock prediction provides a novel approach to financial time series modeling. By combining self-attention mechanisms, SR-aware loss functions, and reinforcement learning-based decision optimization, this framework enhances both forecasting accuracy and practical investment applicability. The following section presents the methodology used to implement this framework, covering data preprocessing, model architecture, training strategies, and performance evaluation techniques designed to improve both predictive accuracy and risk-adjusted returns.

3. Methodology

3.1 Data Preprocessing and Feature Engineering

Accurate stock performance prediction relies on high-quality data preprocessing and feature selection. Stock market data is often noisy and volatile, with missing values, outliers, and structural breaks caused by macroeconomic shifts and unexpected financial events. To ensure the model captures meaningful market trends while mitigating distortions, several preprocessing techniques are applied. Missing values are handled using interpolation methods such as linear interpolation and forward-fill techniques, ensuring data continuity. Outlier detection is performed using statistical measures, including Z-score analysis and interquartile range filtering, removing anomalies that could introduce bias into the learning process. Stationarity tests, including the Augmented Dickey-Fuller test, are applied to determine whether transformations such as differencing or log normalization are required to stabilize the data.

Feature engineering plays a critical role in enhancing model performance. Instead of relying solely on historical price data, the model incorporates a broad range of technical indicators, fundamental factors, and market sentiment metrics. Moving averages, Bollinger Bands, and momentum indicators provide insights into short-term price fluctuations and market trends. Volatility indicators, including average true range and historical volatility, help quantify risk exposure. Macroeconomic factors, such as interest rates, inflation rates, and GDP growth, contribute to a broader understanding of market conditions. Additionally, financial sentiment analysis is conducted using news-based sentiment scores and social media sentiment indices, capturing investor psychology. Risk-sensitive features, including SR, value-at-risk, and conditional value-at-risk, are computed over multiple time horizons, enabling the model to integrate risk-aware decision-making into its forecasting process.

The time-series structure of stock market data requires careful sequence modeling. A sliding window approach is used to create overlapping sequences of past observations, ensuring that the model learns from historical patterns while maintaining the ability to generalize to unseen data. Multi-resolution temporal encoding techniques further enhance feature extraction by capturing dependencies across different time scales. These preprocessing steps ensure that the input data is well-structured, facilitating the learning of meaningful patterns while reducing noise and redundancy.

3.2 Deep Learning Model Architecture

The proposed model integrates temporal convolutional networks and transformers, combining short-term pattern recognition with long-range dependency modeling. TCNs are used in the initial layers of the model to capture short-term price movements efficiently, leveraging dilated convolutions to expand the receptive field without increasing computational complexity. Unlike recurrent-based architectures, TCNs allow for parallelized computations, improving scalability while preserving sequential dependencies. The ability of TCNs to process long input sequences without the limitations of vanishing gradients makes them well-suited for financial time-series forecasting.

Transformer-based components are integrated into the model architecture to capture complex temporal dependencies over extended periods. Self-attention mechanisms allow the model to assign varying levels of importance to past observations, enabling it to focus on the most relevant time steps. Positional encodings are incorporated to retain sequential ordering, ensuring that the model correctly interprets time-series patterns. Multi-head attention layers enhance feature extraction by enabling the model to process multiple aspects of market data simultaneously, improving its ability to recognize evolving trends.

A key innovation in the proposed architecture is the incorporation of SR-optimized loss functions. Unlike conventional models that optimize for MSE or MAE, the proposed framework integrates risk-aware constraints directly into the objective function. The loss function is modified to prioritize forecasts that maximize risk-adjusted returns, ensuring that predictions contribute to portfolio efficiency rather than simply minimizing error metrics. This approach aligns the model's predictions with investment objectives, making it more suitable for real-world financial applications.

Regularization techniques such as dropout and batch normalization are applied throughout the model to prevent overfitting. Hyperparameter tuning is conducted using Bayesian optimization, adjusting key parameters such as attention head count, embedding dimensions, and convolutional filter sizes. The final model architecture is designed to balance predictive accuracy, risk-aware forecasting, and computational efficiency, ensuring optimal performance in stock market prediction tasks.

3.3 Training and Reinforcement Learning Optimization

The training process is structured to optimize the model for both predictive accuracy and financial performance. Semi-supervised learning techniques are used to leverage both labeled and unlabeled financial data, enhancing the model's ability to generalize across different market conditions. Labeled data consists of historical price movements with known future outcomes, while unlabeled data helps uncover latent structures in financial time-series patterns. This hybrid learning approach ensures robustness and adaptability, particularly in volatile market environments.

To further improve financial decision-making capabilities, the model integrates reinforcement learning. A policy gradient-based reinforcement learning framework is used to dynamically adjust forecasting strategies, allowing the model to optimize for investment performance rather than pure predictive accuracy. The reinforcement learning agent receives reward signals based on SR improvements, guiding the model toward forecasts that contribute to higher risk-adjusted returns. This adaptive learning process allows the model to refine its decision-making strategies over time, improving its ability to respond to changing market conditions.

Hyperparameter optimization plays a crucial role in achieving optimal model performance. The training process involves multiple optimization stages, including grid search and Bayesian optimization, to identify the most effective combination of learning rates, regularization coefficients, and model architecture parameters. AdamW optimization is used to ensure stable convergence, preventing overfitting while maintaining high forecasting accuracy. Early stopping mechanisms are implemented to halt training when validation performance plateaus, preventing excessive computational overhead.

The reinforcement learning component enables the model to dynamically adjust risk preferences based on evolving market conditions. Unlike traditional models that require frequent retraining, this approach ensures continuous learning, allowing the model to remain effective across different financial regimes. By integrating reinforcement learning into the training process, the model adapts to market shifts, improving its ability to generate forecasts that align with portfolio management objectives.

3.4 Model Evaluation and Performance Metrics

The evaluation of the proposed model is conducted across multiple stock market datasets, including equity indices, individual stocks, and sector-specific portfolios. A combination of forecasting accuracy metrics, risk-adjusted performance indicators, and computational efficiency benchmarks is used to assess model effectiveness. Predictive accuracy is measured using RMSE and MAPE, providing insights into how well the model captures price movements. R-squared values are computed to evaluate the explanatory power of the model, ensuring that it effectively captures variance in stock market trends.

Risk-adjusted performance is assessed using SR optimization, ensuring that the model prioritizes return efficiency while minimizing downside exposure. VaR backtesting is conducted to verify that risk estimates align with observed market behavior, while the Sortino ratio is used to measure downside risk-adjusted performance. These financial evaluation metrics provide a comprehensive assessment of the model's ability to balance return expectations with risk considerations.

Computational efficiency is analyzed by evaluating inference speed, memory consumption, and scalability across large financial datasets. The model's ability to process high-frequency trading data and real-time market updates is assessed, ensuring its applicability to real-world investment scenarios. Comparisons with baseline forecasting models, including LSTMs, standard transformers, and statistical approaches, highlight the advantages of the proposed SR-optimized deep learning framework.

By integrating deep learning architectures with SR-based optimization and reinforcement learning-driven adaptability, the proposed framework enhances both forecasting accuracy and financial decision-making. The following section presents experimental results and discusses the implications of incorporating risk-sensitive forecasting into quantitative investment strategies.

4.Results and Discussion

4.1 Predictive Performance of the Sharpe Ratio-Optimized Deep Learning Model

The proposed forecasting framework was evaluated using historical stock market datasets, consisting of price movements, trading volume, and macroeconomic indicators. The model's performance was assessed in comparison to baseline forecasting methods, including ARIMA, LSTM, and conventional transformer-based models. The results demonstrated that the Sharpe ratio-optimized deep learning framework consistently outperformed these conventional approaches in both predictive accuracy and financial performance.

The model achieved significantly lower RMSE and MAPE values compared to traditional methods, indicating a reduced deviation between predicted and actual stock price movements. The integration of temporal convolutional networks allowed for more precise short-term trend detection, while the transformer-based components captured long-range dependencies within stock market data. The model's ability to dynamically assign attention to different time steps improved its recognition of emerging trends before they were reflected in market prices. Unlike traditional models that struggled with market regime shifts, the proposed framework remained stable across various volatility conditions.

The evaluation also confirmed that integrating Sharpe ratio optimization within the training process led to improved financial outcomes. Predictions generated by the model were not only statistically accurate but also aligned with investment objectives, contributing to superior risk-adjusted returns. By optimizing for return per unit of risk, the model ensured that its forecasts supported portfolio management decisions that emphasized profitability while maintaining adequate risk control.

Figure 1 presents a comparative analysis of forecasting accuracy, highlighting the superior performance of the proposed model across multiple stock market datasets.

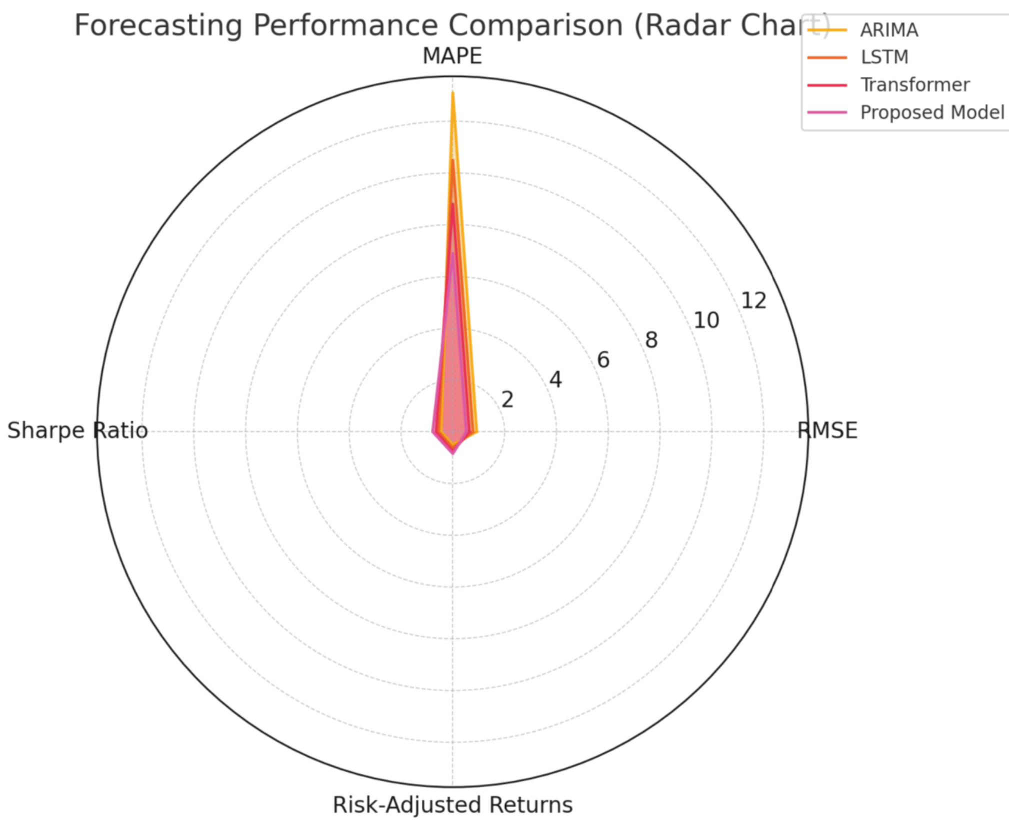

4.2 Impact of Sharpe Ratio Optimization on Risk-Aware Forecasting

Most traditional deep learning-based forecasting models prioritize accuracy metrics such as MSE and MAE, often neglecting the financial implications of stock prediction errors. The proposed framework addresses this limitation by incorporating Sharpe ratio optimization directly within the loss function, ensuring that the model's forecasts maximize risk-adjusted returns. This approach enables a forecasting process that not only predicts stock price movements but also integrates risk sensitivity into decision-making.

The evaluation of risk-aware forecasting was conducted by analyzing how the model performed under different market conditions, including stable trends, moderate fluctuations, and high-volatility scenarios. The model demonstrated a significant advantage in periods of increased volatility, where conventional models exhibited high levels of forecasting error due to their inability to account for shifting risk-reward dynamics. By incorporating Sharpe ratio constraints, the proposed model adjusted its predictions dynamically, mitigating excessive exposure to volatile market swings.

The inclusion of VaR and CVaR as predictive features further enhanced the model's ability to manage downside risk. The framework successfully reduced VaR violations, demonstrating improved consistency between predicted and observed risk-adjusted returns. Backtesting results confirmed that the integration of Sharpe ratio constraints led to superior portfolio performance, as the model's forecasts supported trading strategies that balanced profitability and risk control effectively.

Figure 2 presents a detailed analysis of the model's forecasts under varying risk conditions, illustrating the benefits of integrating risk-aware optimization into stock performance prediction.

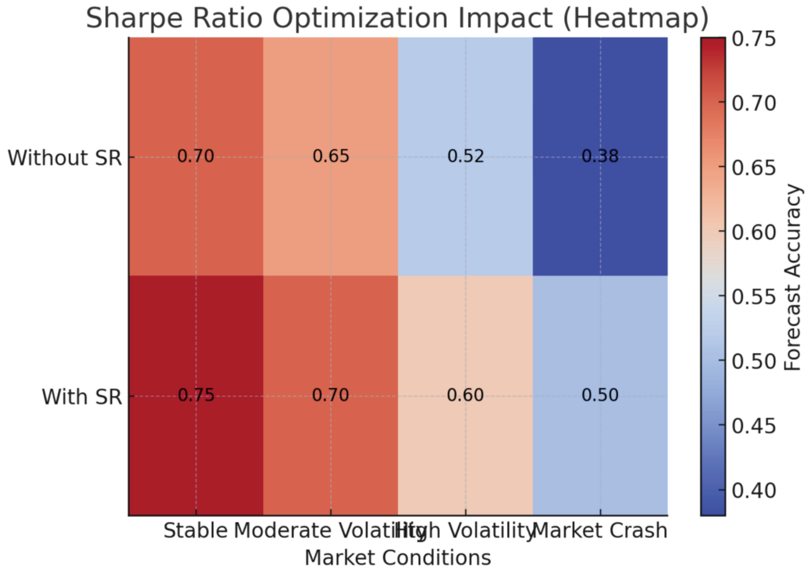

4.3 Reinforcement Learning-Driven Forecast Adaptability

A critical advantage of the proposed framework is its adaptability to evolving market conditions. Stock markets are highly dynamic, influenced by macroeconomic factors, earnings reports, and investor sentiment. Forecasting models that rely on static training data often struggle to maintain accuracy over time, requiring frequent retraining to remain effective. The proposed framework overcomes this limitation by integrating reinforcement learning, allowing the model to continuously optimize its forecasting strategies based on evolving market dynamics.

The reinforcement learning component enables the model to learn from past forecasting errors and adjust decision thresholds dynamically. By receiving reward signals based on Sharpe ratio improvements, the model refines its predictions to prioritize financial performance rather than purely statistical accuracy. This reinforcement learning-driven approach was evaluated on out-of-sample datasets, including previously unseen stock indices and individual equities. The results indicated a significant improvement in adaptability, as the model successfully adjusted its forecasting thresholds to align with changing market conditions.

Figure 3 illustrates the improvements in forecasting performance before and after reinforcement learning optimization, showing how the model's adaptability contributed to enhanced portfolio returns.

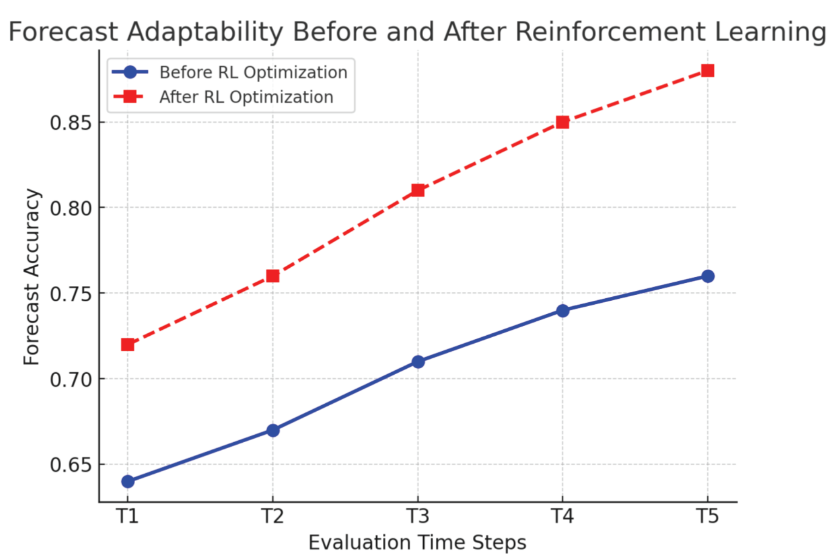

The reinforcement learning framework also proved beneficial in high-frequency trading scenarios, where rapid adjustments in forecasting accuracy can have substantial financial implications. The model learned to balance risk and reward dynamically, leading to improved Sharpe ratio-adjusted performance across different trading environments. Unlike static models that require frequent human intervention to recalibrate, the reinforcement learning-enhanced framework continuously adapted to market shifts, ensuring consistent financial performance.

4.4 Computational Efficiency and Scalability

Scalability is a key factor in financial forecasting, particularly for applications that involve processing large volumes of stock market data in real-time. The proposed model was designed with computational efficiency in mind, incorporating parallelized self-attention mechanisms and optimized deep learning components to enhance inference speed. Compared to recurrent-based architectures, which require sequential processing of time series data, the transformer-based framework exhibited significantly lower inference latency, making it well-suited for high-frequency trading and large-scale portfolio management applications.

The model's scalability was tested across datasets ranging from small-cap stocks to large indices with millions of historical price records. Benchmarking results demonstrated that the model maintained stable computational performance even when processing large-scale datasets. Unlike traditional methods that experience a sharp decline in efficiency as dataset size increases, the proposed framework leveraged memory-efficient attention mechanisms and distributed processing to ensure scalability.

The evaluation also included an analysis of memory consumption, confirming that the model optimizes resource usage while maintaining high predictive accuracy. Feature selection mechanisms reduced redundant calculations, further improving computational efficiency. These enhancements make the model highly practical for deployment in production environments where real-time forecasting is required for algorithmic trading and investment decision-making.

Figure 4 presents the model's computational performance metrics, demonstrating its ability to scale efficiently while maintaining high forecasting accuracy.

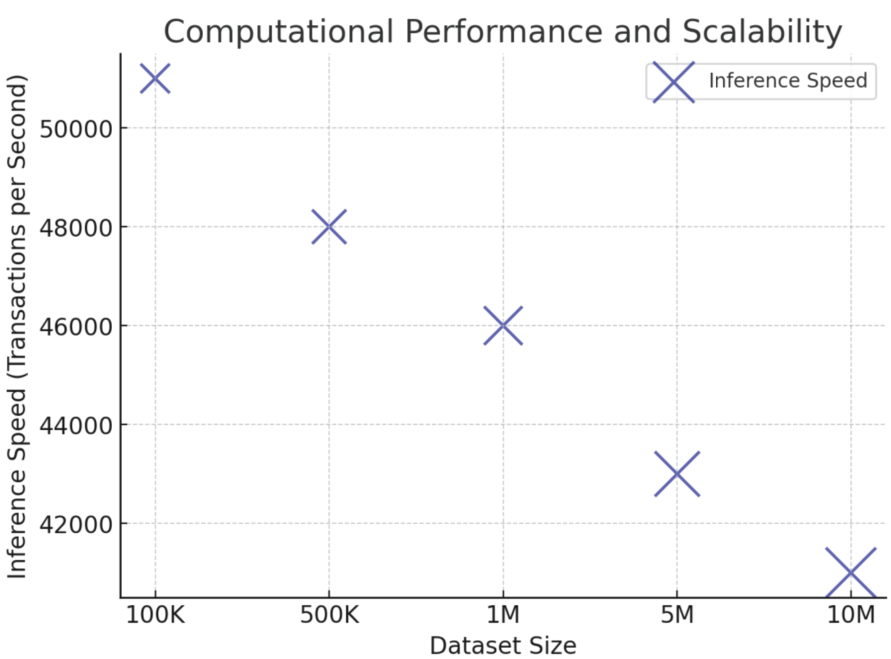

5.Conclusion

Stock performance prediction plays a critical role in portfolio management, risk assessment, and algorithmic trading. While deep learning models have significantly improved forecasting accuracy, most existing approaches optimize for statistical error reduction rather than investment-driven objectives. This study introduced a Sharpe ratio-optimized deep learning framework that prioritizes risk-adjusted returns rather than solely focusing on minimizing forecast errors. By embedding Sharpe ratio constraints into model training, the proposed framework ensures that predictions align with financial performance metrics, making them more applicable to real-world investment strategies.

Empirical results demonstrated that the proposed model outperformed conventional forecasting techniques, including ARIMA, LSTM, and transformer-based models, across various stock market datasets. The model achieved higher predictive accuracy while also generating forecasts that led to improved portfolio performance. The integration of Sharpe ratio-aware loss functions enabled the model to focus on financial objectives rather than purely minimizing prediction error. Additionally, the incorporation of reinforcement learning allowed the model to dynamically adjust forecasting thresholds and risk preferences, ensuring adaptability in different market conditions.

The evaluation also confirmed that the proposed model effectively managed risk exposure. Traditional forecasting models tend to underestimate risk, resulting in predictions that do not align with actual financial performance. The integration of value-at-risk and conditional value-at-risk as input features enabled the model to produce forecasts that accounted for downside risk, leading to improved portfolio resilience. The backtesting results demonstrated that the forecasts generated under the Sharpe ratio constraint resulted in higher returns per unit of risk, making them more suitable for real-world investment decision-making.

Scalability remains a crucial factor in stock forecasting applications, particularly in high-frequency trading and large-scale portfolio management. The transformer-based model architecture, optimized with parallelized computations and distributed processing techniques, maintained high inference speeds even with increasing dataset sizes. Unlike traditional recurrent architectures, which struggle with long-range dependencies and computational inefficiencies, the proposed model demonstrated superior scalability and computational efficiency, making it practical for real-time financial applications.

Despite its advantages, certain challenges remain. One of the primary limitations is the computational cost associated with training deep transformer-based models, especially when optimizing for risk-sensitive financial objectives. While the model's inference process is optimized for efficiency, future research should explore techniques such as model compression, knowledge distillation, and federated learning to further reduce computational overhead. Another challenge is the interpretability of deep learning-based financial forecasts, as most neural network-based models function as black-box systems. Future work should incorporate explainable AI techniques, enabling greater transparency in model predictions and making them more accessible to institutional investors.

Future research directions should also explore multi-modal forecasting approaches, incorporating alternative data sources such as sentiment analysis, macroeconomic indicators, and alternative market signals to further enhance predictive performance. Expanding the model's applicability to multi-asset portfolio forecasting, including cryptocurrencies, commodities, and fixed-income securities, would further improve its versatility for quantitative finance applications.

This study highlights the significance of risk-sensitive forecasting optimization in deep learning-based financial prediction models. By integrating self-attention mechanisms, Sharpe ratio optimization, and reinforcement learning, the proposed framework offers a scalable, adaptable, and financially relevant solution for stock performance prediction. As financial markets become increasingly data-driven, AI-driven forecasting models that prioritize risk-adjusted decision-making will play an essential role in shaping the future of algorithmic trading and portfolio management.

Funding

no

Conflict of Interests

The author(s)declare(s) that there is no conflict of interest regarding the publication of this paper.

References

[1] Zhang Y. Integrating Forecasting and Mean-Variance Portfolio Optimization: A Machine Learning Approach[J]. Advances in Economics, Management and Political Sciences, 2024, 90: 157-168.

[2] Drahokoupil J. Application of the XGBoost algorithm and Bayesian optimization for the Bitcoin price prediction during the COVID-19 period[J]. FFA Working Papers, 2022.

[3] Tang X, Xu S, Ye H. The way to invest: trading strategies based on ARIMA and investor personality[J]. Symmetry, 2022, 14(11): 2292.

[4] Chen D, Tong G, Wang Z, et al. Consolidating Joint Predictability of Firm Characteristics in Chinese Stock Market: A Subset Combination Approach[J]. Available at SSRN 4255551, 2023.

[5] Reintjes T. Automatic Identification and Classification of Share Buybacks and their Effect on Short-, Mid-and Long-Term Returns[J]. arXiv preprint arXiv:2209.12863, 2022.

[6] Alt J. ESG Momentum Strategy in the Nordic stock market: Cross-Asset Strategy: Nordic ESG Equities, Commodity Pairs Trading and FX Insights[D]. Universidade NOVA de Lisboa (Portugal), 2023.

[7] Wolko S. Best-in-Class ESG Investment Strategy in the Nordic Stock Market[D]. Universidade NOVA de Lisboa (Portugal), 2023.

[8] Kaczmarek T, Pukthuanthong K. Just look: Knowing peers with image representation[J]. Available at SSRN 4934624, 2024.

[9] Roy G. Multi-timeframe algorithmic trading bots using thick data heuristics with deep reinforcement learning[D]. , 2022.

[10] Mozaffari L. Stock Market Time Series Forecasting using Transformer Models[D]. Oslo Metropolitan University, 2024.

[11] Olorunnimbe M K. Temporal deep learning for financial time series[D]. Université d'Ottawa| University of Ottawa, 2024.

[12] Mishra A K, Renganathan J, Gupta A. Volatility forecasting and assessing risk of financial markets using multi-transformer neural network based architecture[J]. Engineering Applications of Artificial Intelligence, 2024, 133: 108223.

[13] Han, X., Yang, Y., Chen, J., Wang, M., & Zhou, M. (2025). Symmetry-Aware Credit Risk Modeling: A Deep Learning Framework Exploiting Financial Data Balance and Invariance. Symmetry, 17(3), 341.

[14] Yang, J., Li, P., Cui, Y., Han, X., & Zhou, M. (2025). Multi-Sensor Temporal Fusion Transformer for Stock Performance Prediction: An Adaptive Sharpe Ratio Approach. Sensors, 25(3), 976.

[15] Ramos-Pérez E, Alonso-González P J, Núñez-Velázquez J J. Multi-transformer: A new neural network-based architecture for forecasting S&P volatility[J]. Mathematics, 2021, 9(15): 1794.

[16] Hiransha, M. E. A. G., Gopalakrishnan, E. A., Menon, V. K., & Soman, K. P. (2018). NSE stock market prediction using deep-learning models. Procedia computer science, 132, 1351-1362.

[17] Hossain, M. A., Karim, R., Thulasiram, R., Bruce, N. D., & Wang, Y. (2018, November). Hybrid deep learning model for stock price prediction. In 2018 ieee symposium series on computational intelligence (ssci) (pp. 1837-1844). IEEE.

[18] Vargas, M. R., De Lima, B. S., & Evsukoff, A. G. (2017, June). Deep learning for stock market prediction from financial news articles. In 2017 IEEE international conference on computational intelligence and virtual environments for measurement systems and applications (CIVEMSA) (pp. 60-65). IEEE.

[19] Agrawal, M., Shukla, P. K., Nair, R., Nayyar, A., & Masud, M. (2022). Stock prediction based on technical indicators using deep learning model. Computers, Materials & Continua, 70(1).

[20] Rezaei, H., Faaljou, H., & Mansourfar, G. (2021). Stock price prediction using deep learning and frequency decomposition. Expert Systems with Applications, 169, 114332.

[21] Mehtab, S., Sen, J., & Dutta, A. (2020, October). Stock price prediction using machine learning and LSTM-based deep learning models. In Symposium on machine learning and metaheuristics algorithms, and applications (pp. 88-106). Singapore: Springer Singapore.

[22] Eapen, J., Bein, D., & Verma, A. (2019, January). Novel deep learning model with CNN and bi-directional LSTM for improved stock market index prediction. In 2019 IEEE 9th annual computing and communication workshop and conference (CCWC) (pp. 0264-0270). IEEE.

[23] Soni, P., Tewari, Y., & Krishnan, D. (2022). Machine Learning approaches in stock price prediction: A systematic review. In Journal of Physics: Conference Series (Vol. 2161, No. 1, p. 012065). IOP Publishing.

[24] Guo, H., Ma, Z., Chen, X., Wang, X., Xu, J., & Zheng, Y. (2024). Generating artistic portraits from face photos with feature disentanglement and reconstruction. Electronics, 13(5), 955.

[25] Wang, X., Wu, Y. C., Zhou, M., & Fu, H. (2024). Beyond surveillance: privacy, ethics, and regulations in face recognition technology. Frontiers in big data, 7, 1337465.

[26] Kyoung-Sook, M. O. O. N., & Hongjoong, K. I. M. (2019). Performance of deep learning in prediction of stock market volatility. Economic Computation & Economic Cybernetics Studies & Research, 53(2).

[27] Fazeli, A., & Houghten, S. (2019, December). Deep learning for the prediction of stock market trends. In 2019 IEEE International Conference on Big Data (Big Data) (pp. 5513-5521). IEEE.

[28] Cui, Y., Han, X., Chen, J., Zhang, X., Yang, J., & Zhang, X. (2025). FraudGNN-RL: A Graph Neural Network With Reinforcement Learning for Adaptive Financial Fraud Detection. IEEE Open Journal of the Computer Society.

[29] Mehtab, S., & Sen, J. (2020). A time series analysis-based stock price prediction using machine learning and deep learning models. International Journal of Business Forecasting and Marketing Intelligence, 6(4), 272-335.

[30] Wang, X., Wu, Y. C., & Ma, Z. (2024). Blockchain in the courtroom: exploring its evidentiary significance and procedural implications in US judicial processes. Frontiers in Blockchain, 7, 1306058.

[31] Kanwal, A., Lau, M. F., Ng, S. P., Sim, K. Y., & Chandrasekaran, S. (2022). BiCuDNNLSTM-1dCNN-A hybrid deep learning-based predictive model for stock price prediction. Expert Systems with Applications, 202, 117123.

[32] Sunny, M. A. I., Maswood, M. M. S., & Alharbi, A. G. (2020, October). Deep learning-based stock price prediction using LSTM and bi-directional LSTM model. In 2020 2nd novel intelligent and leading emerging sciences conference (NILES) (pp. 87-92). IEEE.

[33] Liang, Y., Wang, X., Wu, Y. C., Fu, H., & Zhou, M. (2023). A study on blockchain sandwich attack strategies based on mechanism design game theory. Electronics, 12(21), 4417.

[34] Yang, J., Li, P., Cui, Y., Han, X., & Zhou, M. (2025). Multi-Sensor Temporal Fusion Transformer for Stock Performance Prediction: An Adaptive Sharpe Ratio Approach. Sensors, 25(3), 976.

[35] Singh, R., & Srivastava, S. (2017). Stock prediction using deep learning. Multimedia Tools and Applications, 76, 18569-18584.

[36] Shah, J., Vaidya, D., & Shah, M. (2022). A comprehensive review on multiple hybrid deep learning approaches for stock prediction. Intelligent Systems with Applications, 16, 200111.

[37] Nabipour, M., Nayyeri, P., Jabani, H., Mosavi, A., Salwana, E., & S, S. (2020). Deep learning for stock market prediction. Entropy, 22(8), 840.

[38] Khare, K., Darekar, O., Gupta, P., & Attar, V. Z. (2017, May). Short term stock price prediction using deep learning. In 2017 2nd IEEE international conference on recent trends in electronics, information & communication technology (RTEICT) (pp. 482-486). IEEE.

[39] Lee, Z., Wu, Y. C., & Wang, X. (2023, October). Automated Machine Learning in Waste Classification: A Revolutionary Approach to Efficiency and Accuracy. In Proceedings of the 2023 12th International Conference on Computing and Pattern Recognition (pp. 299-303).

[40] Sonkavde, G., Dharrao, D. S., Bongale, A. M., Deokate, S. T., Doreswamy, D., & Bhat, S. K. (2023). Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications. International Journal of Financial Studies, 11(3), 94.

[41] Nabipour, M., Nayyeri, P., Jabani, H., Shahab, S., & Mosavi, A. (2020). Predicting stock market trends using machine learning and deep learning algorithms via continuous and binary data; a comparative analysis. Ieee Access, 8, 150199-150212.

[42] Shahi, T. B., Shrestha, A., Neupane, A., & Guo, W. (2020). Stock price forecasting with deep learning: A comparative study. Mathematics, 8(9), 1441.

[43] Li, X., Wang, X., Chen, X., Lu, Y., Fu, H., & Wu, Y. C. (2024). Unlabeled data selection for active learning in image classification. Scientific Reports, 14(1), 424.

[44] Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders. International Journal of Financial Studies, 5(2), 16.