

An Adaptive Sharpe Ratio-Based Temporal Fusion Transformer for Financial Forecasting

Haoran Yu¹, Xiaodong Feng², Jianhong Liu³

¹Southeast University, Nanjing, China

Abstract

Financial time series forecasting plays a critical role in investment decision-making, risk management, and portfolio optimization. Traditional forecasting models, including statistical approaches and deep learning-based architectures, often focus on maximizing predictive accuracy but fail to incorporate dynamic risk-adjusted performance metrics. The Sharpe ratio, a widely used measure of risk-adjusted return, is typically applied post-forecasting, limiting its potential to guide predictive models during training. To address this limitation, this study proposes an adaptive Sharpe ratio-based temporal fusion transformer (AS-TFT) that integrates risk-aware forecasting mechanisms to optimize financial predictions while considering return volatility.

Experimental evaluations on stock indices, cryptocurrency prices, and commodity markets demonstrate that the AS-TFT outperforms conventional forecasting methods in terms of both predictive accuracy and portfolio returns. The results highlight the importance of integrating risk-adjusted financial performance metrics within deep learning-based forecasting architectures, offering a practical and scalable solution for financial decision-making.

Keywords

Sharpe Ratio, Temporal Fusion Transformer, Financial Forecasting, Risk-Adjusted Returns, Deep Learning, Adaptive Optimization.

1. Introduction

Financial forecasting plays a fundamental role in asset management, quantitative trading, and economic decision-making. Investors and financial analysts rely on accurate predictions to manage risk, optimize portfolios, and develop trading strategies [1]. However, forecasting financial time series is inherently challenging due to the presence of noise, market volatility, and regime shifts caused by macroeconomic and geopolitical events. Traditional statistical models have been widely used for time series analysis, but they often assume linearity and stationarity, making them inadequate for capturing the complex, nonlinear dependencies observed in financial markets. Recent advancements in deep learning have introduced more sophisticated forecasting techniques, with transformer-based models demonstrating state-of-the-art performance [2]. However, existing models largely focus on improving prediction accuracy without explicitly incorporating risk-adjusted performance measures, which are crucial for real-world financial applications.

Traditional forecasting approaches such as autoregressive integrated moving average (ARIMA) and generalized autoregressive conditional heteroskedasticity (GARCH) have been effective in modeling financial data under stable conditions [3]. ARIMA models rely on past observations to predict future values, assuming that market trends follow a deterministic pattern. GARCH models extend this approach by capturing time-varying volatility, making them useful for risk estimation. Despite their effectiveness in some cases, these models struggle to adapt to sudden market shifts and nonlinear dependencies, limiting their predictive power in dynamic financial environments[4].

The emergence of deep learning has revolutionized time series forecasting by enabling models to capture complex temporal patterns and dependencies [5]. Recurrent neural networks

(RNNs) and long short-term memory (LSTM) networks introduced memory mechanisms that improved forecasting accuracy by retaining long-term dependencies[6]. However, these models suffer from vanishing gradient problems, making them inefficient for processing long time series [7]. The introduction of transformer-based architectures overcame these limitations by leveraging self-attention mechanisms, allowing models to process entire time series simultaneously. Transformers have been widely adopted in financial forecasting, demonstrating superior performance in predicting stock prices, volatility, and trading volume [8].

One of the most advanced transformer-based models for time series forecasting is the temporal fusion transformer (TFT). TFT enhances financial forecasting by combining self-attention with recurrent components, enabling multi-horizon predictions. This model effectively learns from both historical trends and future covariates, improving its adaptability to dynamic market conditions[9]. However, despite its advantages, TFT-based models primarily optimize for mean squared error (MSE) or other standard loss functions that focus solely on predictive accuracy. In financial applications, accuracy alone is insufficient; investors require risk-adjusted performance metrics to make informed decisions.

The Sharpe ratio, a widely used measure of risk-adjusted return, is traditionally applied as a post-processing evaluation metric rather than being embedded within forecasting models [10]. The Sharpe ratio quantifies the return per unit of risk, making it an essential criterion for portfolio optimization. However, current deep learning-based forecasting models do not incorporate Sharpe ratio optimization during training, leading to forecasts that may not align with financial objectives. To address this gap, this study proposes an adaptive Sharpe ratio-based temporal fusion transformer (AS-TFT), which integrates risk-adjusted forecasting into the learning process. By embedding an adaptive Sharpe ratio constraint within the model's objective function, the proposed framework ensures that predictions optimize for both accuracy and risk-adjusted returns [11-13].

The AS-TFT framework incorporates reinforcement learning mechanisms to dynamically adjust forecasting thresholds based on evolving market conditions. Unlike static models that require frequent retraining to adapt to market shifts, reinforcement learning enables the model to continuously optimize its decision-making process. Additionally, AS-TFT utilizes multi-resolution temporal encoding, allowing it to capture both short-term fluctuations and long-term investment trends. By integrating risk-sensitive features into transformer-based forecasting, this approach enhances model interpretability and practical applicability in portfolio management[14].

To evaluate the effectiveness of AS-TFT, experiments are conducted on financial datasets including stock market indices, cryptocurrency price movements, and commodity markets. The model's performance is benchmarked against traditional forecasting approaches, recurrent models, and standard transformer implementations. Evaluation metrics include predictive accuracy, risk-adjusted return performance, and portfolio optimization outcomes. The results demonstrate that AS-TFT achieves superior forecasting accuracy while improving Sharpe ratio-based investment returns, highlighting the significance of incorporating risk-aware optimization in deep learning-based financial forecasting.

2. Literature Review

Financial time series forecasting has been extensively studied in both academic research and financial industry applications. Traditional models have primarily relied on statistical methods, while the recent emergence of deep learning-based approaches has significantly improved predictive performance. Despite these advancements, existing forecasting techniques largely focus on minimizing prediction error, neglecting the integration of risk-adjusted performance

metrics that 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 [16].

Early financial forecasting methods were built on statistical time series models that assumed linear dependencies between historical and future values. ARIMA and its variants were widely used due to their ability to model stationary trends, while GARCH models extended this capability by incorporating time-varying volatility estimation. These approaches provided useful insights for financial risk modeling, but their reliance on strict assumptions limited their effectiveness in capturing the nonlinear nature of financial markets. Markets often experience abrupt changes driven by macroeconomic shifts, investor sentiment, and external shocks, making it difficult for these models to adapt to evolving conditions [17-20].

The introduction of machine learning improved financial forecasting by allowing models to learn complex, nonlinear relationships from data. Techniques such as support vector machines and random forests demonstrated enhanced predictive accuracy compared to traditional statistical approaches[21-23]. However, these models did not fully address the sequential dependencies present in financial time series, as they treated observations as independent data points rather than as part of a continuous sequence. To overcome this limitation, RNNs and LSTMs were introduced, offering improved sequential modeling by incorporating memory-based learning. While these architectures performed well in capturing short- and medium-term dependencies, their reliance on sequential processing led to inefficiencies, particularly when dealing with long-range dependencies in high-frequency trading environments [24-27].

Transformer-based models have emerged as a superior alternative, overcoming the scalability limitations of recurrent architectures. Unlike RNNs, transformers process entire time series in parallel using self-attention mechanisms, making them particularly effective for long-range dependency modeling [28-30]. TFT further refined transformer-based time series forecasting by integrating multi-horizon prediction capabilities and feature selection mechanisms. Studies have shown that TFT outperforms both traditional deep learning architectures and conventional forecasting methods in financial applications, achieving superior accuracy in stock price prediction, volatility forecasting, and macroeconomic trend analysis[8]. However, while TFT has demonstrated state-of-the-art performance in predictive accuracy, its training objectives remain primarily focused on minimizing error metrics such as MSE, rather than optimizing for risk-adjusted returns, which are essential in financial applications [31].

Risk-adjusted metrics such as the Sharpe ratio play a crucial role in investment decision-making, providing a measure of return per unit of risk [32]. Despite their importance, existing deep learning-based forecasting models do not integrate these risk-sensitive measures into their optimization frameworks. Instead, the Sharpe ratio is typically used as a post-processing evaluation metric, limiting its impact on the forecasting model's learning process. Recent studies have explored ways to incorporate financial risk metrics into deep learning architectures, demonstrating that embedding risk-sensitive features can improve both predictive robustness and real-world financial applicability [33-35]. However, most of these approaches remain limited to using risk constraints as external components rather than fully integrating them into model training.

The proposed AS-TFT framework directly addresses this limitation by embedding Sharpe ratio-based optimization into the training objective [36-38]. Unlike conventional forecasting techniques that treat price prediction independently from investment risk, AS-TFT ensures that model outputs align with portfolio performance objectives. This is achieved through adaptive risk constraints, which dynamically adjust model forecasts based on changing market conditions[9]. Reinforcement learning is further integrated to optimize the model's forecasting threshold, allowing it to adapt to market regime shifts and volatility fluctuations in real time.

Integrating risk-sensitive optimization into transformer-based forecasting offers a novel approach to financial time series modeling. By combining long-range temporal feature extraction with Sharpe ratio-aware loss functions, AS-TFT enhances both forecasting accuracy and financial applicability. The following section details the methodology used to implement this framework, covering data preprocessing, model architecture, training strategies, and performance evaluation techniques designed to improve both prediction accuracy and risk-adjusted returns.

3. Methodology

3.1. Data Preprocessing and Feature Engineering

Financial time series forecasting requires careful data preprocessing to ensure the accuracy and stability of input data. Raw financial data, including stock prices, exchange rates, and commodity prices, often contain missing values, outliers, and non-stationary patterns that can impact model performance. Missing values are handled using interpolation techniques such as linear interpolation and forward-fill methods to maintain continuity. To detect and correct outliers, statistical measures such as Z-score analysis and interquartile range filtering are applied. Ensuring stationarity in financial data is critical, so transformations like differencing and log normalization are used when necessary.

Feature engineering plays an essential role in enhancing forecasting accuracy. Apart from historical price data, the model incorporates various technical indicators, volatility measures, and macroeconomic factors to improve predictive capability. Momentum indicators such as moving averages and relative strength index capture market trends, while volatility metrics help quantify risk exposure. Adaptive risk measures such as the Sharpe ratio, value-at-risk, and conditional value-at-risk are also computed at multiple time horizons and included as features. These risk-sensitive inputs enable the model to optimize forecasts not only for accuracy but also for financial performance.

A rolling window approach is applied to create sequences of input data, allowing the model to learn from both short-term fluctuations and long-term trends. Multi-resolution temporal encoding further enhances feature extraction by capturing dependencies across different time scales. These preprocessing techniques ensure that the input data is well-structured and that the model can generalize effectively across different financial instruments.

3.2. Adaptive Sharpe Ratio-Based Temporal Fusion Transformer Architecture

The proposed forecasting framework extends the standard temporal fusion transformer by incorporating adaptive Sharpe ratio-based optimization, ensuring that predictions maximize risk-adjusted returns rather than simply minimizing forecast error. The architecture consists of multiple transformer encoder layers with multi-head self-attention mechanisms, enabling the model to capture long-range dependencies within financial time series. By leveraging self-attention, the model can dynamically assign importance to different time steps, improving its ability to recognize significant price movements and emerging market trends.

Gated residual networks are used to enhance feature selection, allowing the model to focus on the most relevant inputs based on prevailing market conditions. Static and time-varying covariate encoders process both fixed financial characteristics and dynamic macroeconomic variables, ensuring that the model captures both asset-specific and broader economic trends. Multi-horizon forecasting capabilities enable the model to generate predictions beyond a single time step, making it well-suited for portfolio management and trading strategies.

A key innovation in this architecture is the integration of Sharpe ratio optimization within the training objective. Traditional forecasting models focus on minimizing absolute or squared errors, which can lead to predictions that do not align with investment goals. By incorporating

a loss function that prioritizes higher Sharpe ratios, the model explicitly learns to balance return expectations with risk management. This approach ensures that forecasted price trends contribute to improved portfolio performance rather than just minimizing deviation from observed values.

Regularization techniques such as dropout and layer normalization are applied throughout the model to improve generalization and prevent overfitting. Positional encodings help retain temporal relationships between input data points, ensuring that the model accurately interprets sequential dependencies within financial time series. By combining these architectural enhancements, the model achieves superior predictive performance while maintaining a strong focus on financial applicability.

3.3. Training and Reinforcement Learning Optimization

The model is trained using semi-supervised learning, allowing it to leverage both labeled and unlabeled financial data. Supervised learning is applied to historical price movements with known outcomes, while unsupervised learning helps uncover latent structures in market behavior. By combining these approaches, the model can generalize effectively even when labeled data is limited.

To ensure that predictions align with risk-adjusted return optimization, reinforcement learning techniques are incorporated into the training process. The model receives reward signals based on Sharpe ratio improvements, guiding it to favor forecasts that contribute to stronger portfolio performance. This reinforcement learning component continuously refines the model's ability to adapt to shifting market conditions, making it highly effective for real-world financial applications.

Hyperparameter tuning is conducted using Bayesian optimization to determine optimal values for attention heads, dropout rates, and learning rate schedules. The model is trained using AdamW optimization, a variant of the Adam algorithm that improves weight decay handling. Early stopping mechanisms are implemented to prevent overfitting by monitoring validation performance and halting training when improvements plateau.

By integrating reinforcement learning into the training process, the model dynamically adjusts its forecasting strategies based on evolving risk-return trade-offs. Unlike static forecasting models that require frequent retraining, this approach enables continuous learning, ensuring that the model remains effective across different market environments.

3.4. Model Evaluation and Performance Metrics

The model is evaluated across multiple financial datasets, including stock indices, cryptocurrency markets, and foreign exchange rates. A combination of traditional forecast accuracy metrics, risk-adjusted performance indicators, and computational efficiency benchmarks is used to assess model effectiveness.

Predictive accuracy is measured using root mean squared error and mean absolute percentage error, providing a quantitative assessment of the model's ability to track actual market trends. R-squared values further evaluate how well the model explains variance in financial time series. Since predictive accuracy alone does not ensure financial usability, additional risk-adjusted performance metrics are incorporated into the evaluation process. Sharpe ratio optimization is used to verify that the model prioritizes return per unit of risk, ensuring alignment with investment objectives. VaR backtesting assesses whether forecasted risk levels are consistent with observed market behavior, while the Sortino ratio provides a downside risk-adjusted measure of performance.

Computational efficiency is another critical evaluation criterion. The model's inference speed is measured to determine its ability to process high-frequency financial data in real time. Memory consumption and scalability tests are conducted to evaluate the model's suitability for

large-scale trading applications. Comparisons with baseline transformer architectures, LSTM-based models, and traditional statistical forecasting techniques highlight the advantages of the proposed AS-TFT framework.

By incorporating risk-aware forecasting strategies and reinforcement learning-based optimization, the model enhances both predictive accuracy and financial applicability. The following section presents experimental results and discusses the impact of integrating risk-sensitive forecasting into deep learning-based financial models.

4. Results and Discussion

4.1. Predictive Performance of the Adaptive Sharpe Ratio-Based Temporal Fusion Transformer

The proposed forecasting framework was evaluated using real-world financial datasets covering a range of asset classes, including equities, foreign exchange, and cryptocurrency markets. To assess predictive accuracy, the model's performance was compared against baseline forecasting methods, including ARIMA, LSTM, and traditional transformer-based models. The results demonstrated that the adaptive Sharpe ratio-based temporal fusion transformer consistently outperformed conventional models in terms of both short-term and long-term predictive accuracy.

The model exhibited significantly lower root mean squared error and mean absolute percentage error compared to baseline approaches. By leveraging multi-horizon forecasting, it effectively captured both short-term price fluctuations and long-term market trends. The incorporation of self-attention mechanisms enabled the model to assign appropriate weights to past observations, improving its ability to detect emerging trends before they materialized in the market. Unlike conventional models that struggled with market regime shifts, the transformer-based framework demonstrated resilience in adapting to different volatility conditions, maintaining high predictive accuracy across diverse market scenarios.

The evaluation also confirmed that the integration of adaptive risk metrics enhanced predictive performance. By incorporating the Sharpe ratio as an optimization criterion, the model generated forecasts that aligned more closely with investment objectives. Predictions were not only accurate in terms of price movements but also contributed to improved portfolio performance by prioritizing returns adjusted for market risk.




Figure 1 illustrates the comparative analysis of forecasting accuracy, highlighting the superior performance of the proposed model across multiple datasets.

4.2. Impact of Adaptive Risk Metrics on Risk-Aware Forecasting

Traditional forecasting models primarily focus on minimizing prediction error while neglecting the financial implications of risk-adjusted returns. The proposed framework addresses this limitation by incorporating adaptive risk metrics, allowing the model to generate forecasts that optimize for both accuracy and financial viability. The evaluation of risk-aware forecasting was conducted by analyzing how the model performed under varying market conditions, including stable trends, moderate volatility, and extreme downturns.

During periods of high volatility, the model demonstrated the ability to adjust its predictions dynamically, mitigating the impact of market fluctuations. Unlike conventional models that tend to overestimate price movements in unstable conditions, the proposed framework maintained stable forecasts by integrating risk-sensitive optimization strategies. The adaptive loss function ensured that forecasts were aligned with risk-adjusted investment objectives, minimizing excessive exposure to high-risk scenarios.

The inclusion of VaR and CVaR as predictive features allowed the model to anticipate downside risk, leading to more balanced forecasting outputs. The model successfully reduced VaR violations, ensuring that risk estimates were consistent with observed market movements. Backtesting results confirmed that the forecasts generated under the adaptive Sharpe ratio constraint led to improved financial decision-making, as they provided better insights into expected returns while maintaining adequate risk control.




Figure 2 presents a detailed analysis of how the model's forecasts adapted to different risk levels, illustrating the effectiveness of integrating risk-aware optimization into the prediction process.

4.3. Reinforcement Learning-Driven Forecast Adaptability

One of the key strengths of the proposed model is its adaptability to evolving market conditions. Financial markets are dynamic, with price movements influenced by macroeconomic events, investor sentiment, and liquidity fluctuations. A static forecasting model would struggle to maintain accuracy over time without frequent retraining, making it impractical for real-world applications. The reinforcement learning component integrated into the proposed framework

allows the model to adjust its forecasting strategies dynamically, improving its ability to respond to sudden market shifts.

The evaluation of adaptability was conducted using out-of-sample datasets containing previously unseen financial instruments. The model demonstrated a significant improvement in prediction accuracy when compared to models without reinforcement learning optimization. It successfully identified changes in market trends and adjusted its predictions accordingly, reducing the lag commonly observed in traditional forecasting models. The reinforcement learning module enabled the model to optimize its decision thresholds continuously, ensuring that forecasts remained relevant even as market dynamics evolved.

The impact of reinforcement learning was particularly evident in high-frequency trading scenarios, where rapid adjustments in forecasting accuracy can lead to substantial differences in portfolio performance. The model learned to balance risk and reward dynamically, leading to an improved Sharpe ratio over extended trading periods.




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

4.4. Computational Efficiency and Scalability

Scalability is a critical factor in financial forecasting, particularly for applications that involve processing large volumes of market data in real time. The proposed model was designed with computational efficiency in mind, utilizing parallelized self-attention mechanisms to improve inference speed. Compared to recurrent-based architectures, which require sequential data processing, the transformer-based model exhibited significantly lower inference latency, making it suitable for high-frequency trading and large-scale portfolio management applications.

The model's scalability was tested across datasets of varying sizes, ranging from 100,000 to 10 million time steps. Benchmarking results showed that the transformer-based framework maintained stable computational efficiency even as dataset sizes increased. Unlike traditional models that experience a sharp decline in performance when handling large datasets, the proposed approach leveraged efficient memory management and distributed processing techniques to ensure scalability. The ability to process large datasets without significant

computational overhead makes the model practical for deployment in production environments where real-time forecasting is required.

The evaluation also included an analysis of memory consumption, which confirmed that the model optimized resource usage while maintaining high predictive performance. By incorporating feature selection mechanisms, the model was able to reduce redundant calculations, further improving its computational efficiency.




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

5. Conclusion

Financial forecasting plays a crucial role in investment strategies, risk management, and portfolio optimization. Traditional models, including statistical and deep learning-based approaches, have primarily focused on minimizing prediction error but have largely neglected the integration of risk-adjusted performance metrics within the forecasting process. The proposed AS-TFT addresses this limitation by incorporating risk-aware optimization techniques, ensuring that financial forecasts align with portfolio management objectives rather than purely predictive accuracy.

Experimental results demonstrated that the proposed model outperforms conventional forecasting techniques, achieving lower RMSE and MAPE values while also improving risk-adjusted returns. By leveraging Sharpe ratio optimization, the model prioritizes forecasts that contribute to higher returns per unit of risk, making it more practical for financial applications. The integration of VaR and CVaR as predictive features allowed the model to anticipate downside risks, leading to more balanced forecasting outputs that minimize excessive exposure to volatility.

A key advantage of AS-TFT is its ability to adapt to changing market conditions. Unlike static forecasting models that require frequent retraining, AS-TFT employs reinforcement learning mechanisms to dynamically optimize forecasting thresholds based on evolving risk-return trade-offs. This adaptability ensures that the model remains effective even as market volatility, macroeconomic conditions, and investor sentiment fluctuate. The experimental evaluation confirmed that the reinforcement learning component contributed to higher Sharpe ratio-

adjusted performance, making the model more suited for real-world investment decision-making.

Scalability remains a crucial factor in financial forecasting applications, particularly in high-frequency trading and large-scale investment portfolios. The proposed model was optimized for computational efficiency, leveraging parallelized self-attention mechanisms and distributed processing techniques to enhance inference speed. The evaluation confirmed that AS-TFT maintains stable computational performance even on large datasets, making it practical for deployment in production environments where real-time forecasting is required.

Despite its advantages, certain challenges remain. One primary limitation is the computational cost associated with training deep transformer models on large-scale financial datasets. While inference speed has been optimized for real-time forecasting, future research should explore model compression techniques and federated learning strategies to further reduce training overhead. Another challenge is interpretability, as deep learning-based financial forecasting models function as black-box systems. Future work should focus on explainable AI techniques, improving transparency in model decision-making to facilitate adoption in institutional finance.

Future research directions should also explore multi-modal forecasting techniques, incorporating additional data sources such as sentiment analysis from financial news and macroeconomic indicators to further enhance predictive performance. Extending AS-TFT to multi-asset forecasting, including cryptocurrency markets, commodities, and fixed-income securities, would also improve its versatility for portfolio managers and institutional investors.

The findings of this study highlight the importance of integrating risk-sensitive optimization within deep learning-based forecasting frameworks. By combining transformer-based modeling, risk-aware forecasting, and reinforcement learning, AS-TFT offers a scalable, adaptable, and financially relevant solution for modern investment strategies. As financial markets continue to evolve, AI-driven forecasting models that align with risk-adjusted decision-making will play a critical role in shaping the future of algorithmic trading and portfolio management.

References

- [1] Päivärinta M. Transformer-based deep learning model for stock return forecasting: Empirical evidence[J]. 2022.
- [2] Xia H, Ao H, Li L, et al. CI-STHPAN: pre-trained attention network for stock selection with channel-independent spatio-temporal hypergraph[C]//Proceedings of the AAAI Conference on Artificial Intelligence. 2024, 38(8): 9187-9195.
- [3] Singh, R., & Srivastava, S. (2017). Stock prediction using deep learning. *Multimedia Tools and Applications*, 76, 18569-18584.
- [4] 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.
- [5] Nabipour, M., Nayyeri, P., Jabani, H., Mosavi, A., Salwana, E., & S, S. (2020). Deep learning for stock market prediction. *Entropy*, 22(8), 840.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] Drahoukoupil J. Application of the XGBoost algorithm and Bayesian optimization for the Bitcoin price prediction during the COVID-19 period[J]. *FFA Working Papers*, 2022.
- [10] Tang X, Xu S, Ye H. The way to invest: trading strategies based on ARIMA and investor personality[J]. *Symmetry*, 2022, 14(11): 2292.

- [11] 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.
- [12] 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.
- [13] 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.
- [14] Wolko S. Best-in-Class ESG Investment Strategy in the Nordic Stock Market[D]. Universidade NOVA de Lisboa (Portugal), 2023.
- [15] Olorunnimbe M K. Temporal deep learning for financial time series[D]. Université d'Ottawa| University of Ottawa, 2024.
- [16] 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.
- [17] 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.
- [18] Wei Z, Chen D, Zhang Y, et al. Deep reinforcement learning portfolio model based on mixture of experts[J]. Applied Intelligence, 2025, 55(5): 347.
- [19] Akşehir Z D, Kılıç E. Analyzing the critical steps in deep learning-based stock forecasting: a literature review[J]. PeerJ Computer Science, 2024, 10: e2312.
- [20] Novykov V, Bilson C, Gepp A, et al. Deep learning applications in investment portfolio management: a systematic literature review[J]. Journal of Accounting Literature, 2023.
- [21] Kodurupaka N, Basavadeepthi H M, Pecheti S T, et al. Deep reinforcement learning in stock trading: Evaluating ddpg and dqn strategies[C]//2024 International Conference on Emerging Smart Computing and Informatics (ESCI). IEEE, 2024: 1-7.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] Liu, Y., Wu, Y. C., Fu, H., Guo, W. Y., & Wang, X. (2023). Digital intervention in improving the outcomes of mental health among LGBTQ+ youth: a systematic review. *Frontiers in psychology*, 14, 1242928.
- [26] 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.
- [27] 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.
- [28] 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.
- [29] 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).
- [30] 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.
- [31] 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.
- [32] 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.
- [33] Noh Y, Kim J M, Hong S, et al. Deep learning model for multivariate high-frequency time-series data: financial market index prediction[J]. Mathematics, 2023, 11(16): 3603.
- [34] Dip Das J. Machine learning and data science application for financial price prediction and portfolio optimization[J]. 2024.
- [35] 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.

- [36] 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).
- [37] das Neves L R. Optimizing Financial Portfolios Through Time: A Comprehensive Machine Learning Approach[D]. Universidade NOVA de Lisboa (Portugal), 2024.
- [38] 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.