

Machine Learning for Binary Classification of Daily Stock Price Movements in the Chinese A-Share Market: A Narrative Review

Qinyuan Wu*

Yunnan University Affiliated Xingyao High School, Yunnan, China

**Corresponding author: Qinyuan Wu.*

Abstract

The stock market exhibits characteristics such as high volatility, nonlinearity, and strong noise. Precisely predicting stock price trends has long been a challenging issue in quantitative finance. Current research often conflates the two prediction tasks: stock trend classification and price return regression. Common issues include inappropriate model selection and unclear applicability to the scenario. In response to the research challenges mentioned, this paper focuses on stock trend prediction, identifies the drawbacks of mainstream models, and provides references and guidance for subsequent research. This article adopts the narrative review method. It examines the empirical literature published from 2020 to 2025, comparing the representative methods such as SVM, XGBoost, LSTM, CNN-LSTM, and Transformer from four aspects: model principles, experimental design, prediction performance, and application scenarios. As a result, different models have their own applicable scenarios based on data characteristics, market environment, and prediction periods. The cost-benefit ratio of XGBoost in the low-frequency prediction scenario is relatively high. LSTM is the current mainstream model in this field. The transformer only has a slight performance advantage in high-frequency and large-sample scenarios. The out-of-sample accuracy of daily stock trend prediction has remained stable between 52% and 60% over the long term. Currently, several common problems in this field include sample category imbalance, non-standard time-series experimental designs that lead to artificially high prediction accuracy, misalignment between classification accuracy and real trading returns, and a lack of interpretability in deep learning models. Based on the research findings, this paper proposes optimization paths across four dimensions: experimental norms, evaluation system, model interpretability, and multimodal feature construction, to advance the intelligent stock trend prediction model towards standardization, practicality, and interpretability.

Keywords

stock trend prediction, machine learning, model comparison, narrative review, financial time series

1. Introduction

In this review, stock trend prediction is defined as a daily binary classification task in the A-share market. This task is based on information available before trading day T , and the model predicts whether the closing price on the next trading day, $T+1$, will rise or fall. The prediction window is one trading day. The empirical studies reviewed in this article have consistently maintained the scope of this task - namely, the A-share market, the daily frequency, and the binary labels of rise or fall. The stock market serves as the primary vehicle for resource allocation in the capital market. Its price fluctuations directly affect investors' returns, market stability, and the macroeconomy. However, stock market conditions are influenced by multiple factors, including macroeconomic conditions, industry cycles, market sentiment, and cross-border capital flows. They exhibit typical characteristics of high volatility, non-linearity, strong noise, and weak stationarity. Accurately identifying short-term trend changes has long been a key challenge in quantitative finance and financial engineering [1]. In recent years, China has vigorously promoted the application of “artificial intelligence +” to the real economy and financial markets. The “Financial Technology Development Plan (2022-2025)” clearly states that by applying big data and artificial intelligence technologies, the capital market's intelligent analysis and risk warning capabilities can be enhanced, and technologies such as quantitative investment research and intelligent risk control can be implemented and applied [2]. Intelligent, quantitative, machine-learning-driven stock prediction technologies have been widely applied in A-share quantitative research and investment scenarios. They effectively enhance the efficiency of trend identification and risk prediction and serve as an important support for intelligent decision-making by institutions. Traditional statistical models such as ARIMA and GARCH are highly dependent on the assumption of linear stationarity. They are ill-suited to the nonlinear, stochastic characteristics of the stock market. They also fail to meet the current requirements for the accuracy and stability of intelligent investment research [3]. Therefore, systematically analyzing the application patterns of various intelligent models for stock trend prediction holds both theoretical and practical significance.

Current research using machine learning for stock prediction has achieved significant results. To establish a clear and consistent classification system, this paper defines the classification axis as a core feature-extraction and representation paradigm, rather than as strictly mutually exclusive architecture types. Based on empirical literature from 2020 to 2025, the main models covered in this paper can be classified into three paradigm families: (1) shallow discriminative models based on manually designed features, represented by SVM and XGBoost, relying on predefined techniques or fundamental indicators; (2) sequence-dependent models based on recurrent neural network architectures, represented by LSTM and CNN-LSTM, capable of automatically capturing short-term and long-term temporal correlations in the price sequence; (3) global interaction models based on attention mechanisms, represented by Transformer, capable of simultaneously learning pairwise dependencies between all time steps. It should be emphasized that this paradigm-based classification focuses on the primary signal-extraction path. For example, the attention module can be embedded as an enhancement term into LSTM, and such hybrid models will be discussed in subsequent chapters according to their dominant paradigm. Other architectures - including GRU, TCN, graph neural networks, and various ensemble models - occasionally appear in related research, but due to their limited application in the A-share daily vector classification research discussed in this paper, they have not been included in the main empirical scope of this review. Existing literature generally confirms that models across these three paradigms outperform traditional statistical methods such as ARIMA and GARCH in nonlinear fitting, but each model exhibits different performance patterns with respect to data scale, feature engineering, and market environment [1,3,4]. However, most existing reviews focus only on enumerating algorithm variants or comparing overall prediction accuracy, conflating directional-trend binary classification with price-return regression. This confusion - especially the failure to strictly distinguish the differences in dataset construction rules, evaluation metrics, and application scenarios between the two types of tasks - leads to inconsistent conclusions and limited practical reference value [5]. At the same time, the existing literature has two other issues: insufficient attention to the adaptability of models to specific market conditions and a lack of comprehensive analysis of common empirical flaws. In conclusion, a systematic review specifically focusing on direction classification tasks and maintaining a clear and consistent model classification system is still significantly lacking to this day.

Based on the aforementioned research background and existing gaps, this paper focuses solely on classifying stock price trends as up or down. This paper systematically reviews the relevant empirical literature from 2020 to 2025, clarifies the iterative development process, performance differences, and applicable

scenarios of each model, and summarizes the common problems and development bottlenecks in current empirical research. To achieve the aforementioned research objectives, this paper systematically presents three core research questions that run through the entire text: “Which machine learning models are most commonly used?” “Which factors affect the performance of different models?” and “What are the common limitations of current research, and what are the future directions?”

The article briefly introduces the basic types of mainstream machine learning and deep learning models for stock prediction and analyzes the technical advantages and application limitations of each model. Then, it provides a detailed description of the architectural features, operational principles, and time-series feature-extraction mechanisms of various time-series prediction models currently in use. Based on this, the paper summarizes empirical results on stock trend prediction directions over the period 2020-2025, examines the core structures, applicable data types, performance indicators, and differentiated application scenarios of various models. It also presents a comprehensive overview of the current state of research, patterns of technical iteration, and core technical bottlenecks in this field.

2. Literature review

2.1 Research Phase of Traditional Machine Learning and Ensemble Learning

In the early stages of stock trend prediction research, traditional machine learning algorithms were the core tools for classifying price movements. Support vector machines (SVMs) were the most widely used. SVM performs binary classification of stock price rise and fall samples by constructing an optimal classification hyperplane. It demonstrates stable classification performance on daily stock data with small sample sizes, low noise, and simple feature dimensions. This makes it suitable for early quantitative research scenarios in which data volume is small and technical indicators are limited [6]. However, SVM also has fundamental flaws in practical applications. In most of the studies reviewed in this paper, the prediction effect relies heavily on manually constructed technical indicators such as MACD, RSI, and moving averages. Although SVM can, in principle, accept input features from raw price windows, autoencoder representations, text or sentiment features, or other learners, such extensions are rarely observed in the reviewed empirical studies. Relying mainly on manual features, it is unable to uncover long-term correlation patterns in time series data, and its generalization ability is poor in the complex, highly random real stock market [4].

With the continuous improvement and maturation of ensemble learning techniques, Extreme Gradient Boosting Tree (XGBoost) has gradually replaced traditional models such as SVM and Random Forest. It has become the mainstream algorithm for daily stock trend prediction due to its comprehensive performance advantages [4,6]. XGBoost includes L1/L2 regularization to mitigate overfitting and a built-in mechanism for handling missing values, which gives it strong robustness against abnormal fluctuations and missing values in stock data. Its tree-based additive structure also supports high training efficiency and low computational cost. These advantages, however, are conditional: gradient boosting does not automatically model time dependence or prevent look-ahead leakage, so in financial series with weak periodicity and strong noise, its performance still depends on carefully constructed lag features, rolling validation, regularization, and stability checks. In the reviewed studies, XGBoost shows a good cost-benefit ratio in regular daily-frequency scenarios [6]. However, XGBoost still suffers from the limitation of relying on manually engineered features and is unable to autonomously discover the deep temporal patterns in the original K-line and trading volume sequences. The model's performance has a fixed upper limit and is difficult to adapt to complex market scenarios characterized by high frequency and intense volatility [4]. In conclusion, traditional machine learning and ensemble learning are based on the core paradigm of manually constructing features. These models are lightweight and highly interpretable, but they have insufficient time-series modeling capabilities and market adaptability, and are only suitable for low-complexity, conventional daily-frequency prediction scenarios.

2.2 The Rise of Sequential Deep Learning

To overcome the technical bottleneck of traditional machine learning, which relies on artificial features, deep learning has been widely applied to financial time-series prediction, transforming stock trend prediction from a human-feature-driven approach to a data-driven one [3]. Traditional recurrent neural networks (RNNs) suffer from vanishing and exploding gradients when training on long sequences, and thus are unable to capture

long-term dependencies in stock time series data. Long Short-Term Memory (LSTM) networks optimize temporal memory through the input, forget, and output gates. They effectively address the inherent defects of RNN. They can be trained directly on the original K-line and trading volume data, autonomously discovering the inherent rise-and-fall patterns, and they do not require humans to construct an indicator system. With its outstanding nonlinear fitting ability and intelligent characteristics, LSTM has become the mainstream model for stock trend direction prediction over the period 2020-2025 [1,3]. CNN-LSTM, as a classic hybrid time-series model, combines the local feature-extraction capability of a Convolutional Neural Network (CNN) with the time-series memory capability of an LSTM. This enables it to simultaneously extract the local features and long-term temporal correlations of stock data. In daily-frequency prediction scenarios with relatively stable trends, some reviewed studies report that CNN-LSTM outperforms LSTM; however, whether this advantage holds depends on the convolution window, feature representation, data volume, and validation protocol, rather than being an inherent property of the architecture. CNN-LSTM is currently the mainstream approach in research on hybrid models [7,8].

2.3 Transformer Model Based on Self-Attention Mechanism

It should be noted that the Transformer is not a single unified architecture: one can distinguish the vanilla Transformer from time-series-oriented variants (such as Informer and PatchTST) and from efficient-attention variants, and design choices such as positional encoding, sequence length, and feature dimension can substantially affect the results. In the empirical studies reviewed in this paper, however, most applications adopt the standard Transformer encoder architecture, and time-series variants with efficient attention remain rare in the A-share daily trend classification literature. With the continuous iteration of deep learning technology, the Transformer model, based on the self-attention mechanism, has overcome the limitations of recurrent neural networks' sequential processing and can compute correlation weights at any time point in a time series. It has significant theoretical advantages for long-time-series modeling and has seen a steady increase in popularity in research on financial prediction in recent years [1]. From a theoretical perspective, the Transformer can precisely capture the hidden evolution patterns in long time-series data and has modeling potential superior to that of traditional time-series models such as LSTM. However, in real-world stock prediction, its theoretical advantages do not fully translate into performance gains.

The stock market exhibits characteristics of strong noise, high randomness, and unstable trends. The Transformer model has a large number of parameters and a complex structure. During training, it is prone to fitting random market noise rather than the actual rise-and-fall patterns. The risk of overfitting is significantly higher than that of XGBoost and LSTM models [1]. Existing empirical studies have shown that in the conventional daily stock trend classification task, the accuracy improvement of Transformer over LSTM is negligible, while the computational power consumption and training costs increase exponentially [3]. The standard Transformer encoder has only a slight performance advantage in ultra-high-frequency trading and in scenarios involving extremely large sample datasets. It cannot replace the dominant position of LSTM and XGBoost in conventional daily-frequency predictions [1]. Overall, the attention mechanism model is at the forefront of financial time series prediction, but it is currently in a stage where “the theory is advanced, but the implementation is limited,” and has not yet developed into a mature, universal industry application model.

3. Comparison of Mainstream Models and Analysis of Applicable Scenarios

A review of the existing empirical literature shows no absolute superiority or inferiority among stock trend prediction models; only differences in scenario adaptation exist. The performance of the models is mainly influenced by factors such as data frequency, sample size, and market conditions [1,3,4]. This chapter conducts a horizontal comparison of the inherent characteristics, core advantages, and inherent limitations of various mainstream models, clarifies their differentiated applicable scenarios, and provides theoretical support for the model selection in stock trend prediction research.

In traditional machine learning models, SVM is suitable for regular daily-frequency prediction scenarios with small sample sizes, low noise, and simple features, and its classification performance is stable, but its generalization ability is weak. XGBoost has strong anti-interference ability, low training cost, and the best comprehensive performance. It is the optimal choice for regular daily-frequency stock predictions with limited computing power and moderate sample size [4]. The common drawback of these two models is that they rely

too heavily on manual feature engineering and cannot autonomously extract complex time-series features. Their performance has a fixed upper limit and is difficult to adapt to high-frequency and highly volatile market environments.

The LSTM and CNN-LSTM models are suitable for trend-feature-dominant daily-frequency data and can autonomously discover temporal evolution patterns without the need for manual indicator construction [3]. Among them, LSTM is generally regarded in the reviewed literature as the most universal, in the sense of maintaining relatively stable performance across different stocks, periods, and feature sets; however, this conclusion is mostly based on within-sample comparisons, and external validation or rolling-window tests across different market states remain limited. LSTM effectively balances model complexity and prediction stability, making it suitable for most conventional daily-frequency prediction scenarios. CNN-LSTM offers richer feature extraction capabilities and can integrate local and temporal features, but it has high data quality requirements. Whether it overfits in volatile markets depends on regularization, data scale, window design, and the validation protocol; in the reviewed studies, it tends to perform better on high-quality time series data with relatively stable trends, but this should not be treated as an architecture-inherent boundary [7,8]. Both models are highly sensitive to hyperparameters, and improper parameter settings can significantly degrade performance.

Although Transformers possess top-notch theoretical modeling capabilities, their practical application has significant limitations. It is suitable only for specific research scenarios involving ultra-high frequencies, large sample sizes, and sufficient computing power [1]. In the routine daily stock prediction task, Transformers cannot achieve substantial improvements in accuracy. Blind use, divorced from data characteristics and task requirements, will only waste computing resources and cannot optimize predictive performance.

In conclusion, when selecting models for stock trend prediction, the principle of “scenario adaptation and cost-effectiveness priority” should be followed. We should abandon the misconception that “the newer the model, the better the effect,” and instead choose the optimal model based on data characteristics, computing power conditions, and research requirements.

4. Summary of Common Phenomena in Existing Empirical Studies

A review of 7 core empirical studies from 2020 to 2025 reveals a clear hierarchical pattern in the application of models for stock trend classification. LSTM, with its outstanding comprehensive adaptability, has the highest application frequency and holds a dominant position in the industry; XGBoost, with its high cost-effectiveness, is widely used in regular daily prediction scenarios; the traditional SVM model, due to its reliance on fixed manual features, its lack of end-to-end time-series representation learning, and its performance decline in some cross-period tests, has seen a decreasing proportion of applications over time; the theoretical research on Transformer has been increasing year by year, but there are very few empirical studies with actual performance gains, and most results only remain at the theoretical comparison level [1,3].

Based on empirical performance metrics such as prediction accuracy, the results of most existing daily-frequency A-share price-up/down binary classification studies are highly consistent. Across the reviewed daily-frequency studies, the reported out-of-sample accuracy of mainstream models has largely clustered within 52%–60%, and substantial gains above this range remain difficult to achieve despite structural optimization and feature iteration [3,5]. Note that this range reflects cross-study observation rather than a single fixed benchmark; [1] in particular reports results on high-frequency limit order book data and is cited here only as a related high-frequency reference. This accuracy range is not an accidental result of a single study, but a common feature of daily-frequency stock trend classification in global capital markets, and has universal applicability [5].

Apart from the fixed accuracy limit, current empirical research exhibits numerous non-standard and inconsistent issues that seriously affect the authenticity and comparability of the results and hinder technological iteration in the field. Firstly, the division of time series datasets is not standardized. Some studies violate the principle of continuous time in financial sequences by randomly splitting the training and test sets, thereby leaking future information, yielding falsely high predicted accuracy and leaving the research results lacking practical value [3]. Secondly, the evaluation standards are not fully consistent. The thresholds for rise and fall returns, as well as the out-of-sample prediction windows, vary across studies. Such differences may

reflect different trading objectives and noise-filtering strategies rather than a lack of rigor, but because thresholds, holding periods, and class distributions are often not fully reported, horizontal cross-study comparisons become difficult [3]. Thirdly, the evaluation system is single-sided. Most studies use overall accuracy as the sole indicator, ignoring key metrics such as F1 score, precision, and recall, which cannot capture prediction bias caused by sample imbalance and make it difficult to objectively measure the model's true performance [9].

5. Core Research Challenges and Reflections

5.1 The Problem of Imbalanced Sample Categories

Stock market trends exhibit distinct cyclical and phased characteristics. During volatile markets, the distribution of rising and falling samples is relatively balanced. However, extreme scenarios, such as a single-sided bull or bear market, can lead to a severe imbalance in the data set's classification. Model training can lead to prediction bias, where the model tends to fit the majority class samples with a higher proportion of occurrences. As a result, the overall accuracy appears inflated, but the model's ability to identify extreme market conditions is very poor. Extreme market conditions are precisely the core nodes of the gains and risks in quantitative trading.

Meanwhile, the financial time series has strict temporal correlations. Traditional oversampling and undersampling methods will disrupt the temporal logic and lead to information leakage. Not only do they fail to address sample imbalance, but they also reduce the model's ability to generalize, becoming a core pain point in the field [4].

5.2 The Classification Accuracy Rate is Disconnected from the Actual Trading Returns

Currently, there is a widespread misconception in the academic community that “prediction accuracy is equivalent to trading returns,” which has led to the failure of many high-performance models in the laboratory to be applied in real trading scenarios. The core reason is an essential mismatch between classification accuracy and actual trading returns [10]. Classification accuracy only determines whether the direction of price movement is correct, completely ignoring differences in the magnitude of price fluctuations. Most models increase their overall accuracy by identifying numerous small, insignificant fluctuations, but they are unable to predict key market conditions that drive significant price movements accurately. In real trading, multiple small profits are easily offset by a single large loss, resulting in the phenomenon of “high backtesting accuracy but persistent real-trading losses.” A single accuracy indicator, divorced from the real trading environment, cannot capture the model's practical profitability and seriously hinders its industrial implementation [10]. More generally, actual trading returns depend on position sizing, the chosen threshold, turnover, transaction fees, slippage, and the distribution of gains and losses. A rigorous evaluation should therefore report net returns, Sharpe/Sortino ratios, maximum drawdown, turnover rate, and statistical significance, rather than relying on raw accuracy alone [10].

5.3 Real-time Trading Scenario Signal Delay and Profit Loss

The existing empirical studies are all conducted in an ideal experimental environment, use the T+1 closing price to construct up/down labels, and ignore various constraints in real trading. In real trading, the model's predicted signals at time T can only be executed at the opening of T+1. Real factors such as overnight news impact, opening gap, liquidity fluctuations, transaction fees, slippage losses, etc., will cause signal lag and significant loss of returns. The neglect of the constraints of the experimental environment in actual trading leads to a serious overestimation of the model's performance, and laboratory backtesting results show a large gap from actual market performance [10].

5.4 Model Interpretability Deficiency and Optimization Path

The mainstream deep learning models, such as LSTM and Transformer, are generally regarded as black box models. Although they can output probabilities, confidence scores, hidden representations, and intermediate attention weights, these outputs cannot automatically constitute stable, reliable, and auditable explanations of the decision basis, core influencing features, and risk sources [11]. The financial market has

strict compliance requirements and high risk sensitivity. Unexplainable prediction signals cannot support institutions in risk management, decision traceability, and compliance verification, thereby significantly limiting the use cases for intelligent prediction models. To address the interpretability issue, the current academic community has formed two main optimization paths: One is post-event interpretability analysis, which relies on SHAP and LIME tools to quantify the contribution weights of each feature and visualize the decision logic of the model; however, these post-event methods, although having high credibility in local explanations, lack stability across different data subsets, and their audibility is still limited. The other is intrinsic interpretability design, which achieves traceability of the prediction process through attention-weight visualization and gating mechanisms, but attention weights do not always faithfully reflect the actual decision logic, so additional verification is required [11]. Moreover, integrating multimodal features such as macroeconomic factors, market opinions, and investor sentiment can effectively compensate for the shortcomings of a single technical indicator, thereby enhancing the credibility and stability of prediction results [8].

6. Conclusion and Future Prospects

This paper focuses on the daily binary classification of stock price trends in the A-share market, strictly distinguishing trend classification from return regression. Based on seven high-quality empirical studies from 2020 to 2025, it reviews the iterative development, performance characteristics, and applicable scenarios of mainstream machine learning and deep learning models, and answers its three research questions. First, the mainstream models are SVM, XGBoost, LSTM, CNN-LSTM, and Transformer, each with its own applicable scenario rather than absolute superiority: SVM suits small-sample, low-complexity daily tasks; XGBoost is the optimal choice for regular daily-frequency predictions with limited computing power; LSTM has the strongest universality for most daily trend markets; CNN-LSTM has outstanding feature extraction but a higher overfitting risk; and Transformer only shows a slight advantage in extremely large-sample, ultra-high-frequency, high-computing-power scenarios. Second, data frequency, prediction window, input features, and market conditions significantly affect model performance, with market conditions as the core factor, while simply iterating the model structure has a limited effect on accuracy. Third, current research has four common limitations: imbalanced sample categories, non-standard time-series experimental design, a single accuracy-based evaluation system that is disconnected from real trading returns, and insufficient model interpretability. Over the past five years, model structures have continuously evolved, but the out-of-sample accuracy of daily trend prediction has remained stable at 52% - 60%, indicating that the core bottleneck is not algorithm performance but rather market noise, experimental design, evaluation systems, and interpretability.

6.1 Future Research Prospects

Taking into account the limitations of existing research and the industry's development trends, this paper proposes four future research directions, each presented as a verifiable question. First, regarding experimental norms: In daily A-share data, which rolling verification scheme (such as forward rolling or extended window division) most effectively controls information leakage, and how sensitive are the reported results to the choice of verification protocol? Future research should not require all studies to adopt the same rise and fall thresholds; instead, it should fully disclose the thresholds, holding periods, and category distributions, and provide sensitivity analyses across multiple thresholds, because the thresholds should match the holding period, transaction costs, and trading goals [1,3]. Second, regarding the evaluation system: Under the assumption of clearly defined thresholds, holding periods, transaction costs, and slippage, how should cost-sensitive evaluation be designed to make the model's ranking under the original accuracy rate consistent with the ranking under net yield, Sharpe/Sortino ratio, maximum drawdown, turnover rate, and statistical significance? Third, regarding interpretability: How to verify the performance of SHAP, LIME, and attention mechanism-based explanations in terms of fidelity, robustness, and reproducibility (for example, through perturbation tests or by comparing with ablation results), and whether these explanations truly enhance risk control and compliance decisions? SHAP, LIME, and attention visualizations themselves cannot solve the black box problem; they only provide approximate explanations and may be unstable [11]. Fourth, regarding multimodal features: How should macro data, market sentiment, and investor sentiment be aligned at the release time to avoid forward-looking bias? Can multimodal fusion surpass the baseline of a single modality and improve the out-of-sample accuracy [8]? To answer these questions, we suggest establishing a clear set of standard task sets, including

fixed stock pools, unified daily binary labels, and forward progression division methods, along with a minimum report list covering thresholds, holding periods, category distributions, and cost assumptions. Future research on stock trend prediction will gradually move away from the misconception of “emphasizing models, ignoring scenarios, emphasizing accuracy, and ignoring practicality” and develop towards standardization, practicality, and interpretability.

References

- [1] Bilokon, P., & Qiu, Y. (2023). Transformers versus LSTMs for Electronic Trading. arXiv preprint arXiv:2309.11400.
- [2] People's Bank of China. (2021). Financial Technology Development Plan (2022-2025) [Policy document]
- [3] Zou, J., Zhao, Q., Jiao, Y., et al. (2022). Stock Market Prediction via Deep Learning Techniques: A Survey. arXiv preprint arXiv:2212.12717.
- [4] Chen, Q. (2022). Stock Market Prediction Using Machine Learning. ICBEM 2022, Atlantis Press, Atlantis Highlights in Intelligent Systems, 5, pp. 458-465.
- [5] Sonkavde, G., Dharrao, D. S., Bongale, A. M., et al. (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.
- [6] Zhang, X., & Wen, X. (2023). Research on Stock Index Rise-Fall Prediction Strategy Based on XGBoost. Computer & Digital Engineering, 51(03), 686-689.
- [7] Zhang, J., Ye, L., & Lai, Y. (2023). Stock Price Prediction Using CNN-BiLSTM-Attention Model. Mathematics, 11(9), 1985.
- [8] Sun, Y., Sun, Q., & Zhu, S. (2022). Prediction of Shanghai Stock Index Based on Investor Sentiment and CNN-LSTM Model. Journal of Systems Science and Information, 10(6), 620-632.
- [9] Ranjbaran, G., Recupero, D. R., Lombardo, G., & Consoli, S. (2023). Leveraging augmentation techniques for tasks with unbalancedness within the financial domain: a two-level ensemble approach. EPJ Data Science, 12, 24.
- [10] Wang, Y., & Yan, K. (2023). Machine learning-based quantitative trading strategies across different time intervals in the American market. Quantitative Finance and Economics, 7(4), 569-594.
- [11] Chen, H., Lundberg, S. M., & Lee, S. I. (2022). Explaining a series of models by propagating Shapley values. Nature Communications, 13, 4512.

Funding

This research received no external funding.

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Open Access

This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

