Research on Stock Price Prediction Based on LSTM Hybrid Model

Tongshu Wu
Dept. of AI and Big Data, Xi'an Jiaotong-Liverpool University. Suzhou, China

Abstract: Stock price forecasting is an important topic in financial time series analysis. Traditional time series models have limitations in processing nonlinear financial data, and deep learning models show application potential due to their strong feature learning capabilities. In this study, a long short-term memory network and an autoregressive comprehensive moving average model are constructed to systematically compare their performance differences in stock price prediction. The study adopts a unified technical index system and a multi-dimensional evaluation framework to evaluate the performance of the model from three dimensions: prediction accuracy, direction judgment and generalization ability. The empirical results show that the deep learning model has significant advantages in capturing the nonlinear characteristics of stock prices, and is better than the traditional time series method in all evaluation indicators. It is manifested in higher prediction accuracy, better trend judgment ability, and stronger out-of-sample generalization performance. The research results confirm the effectiveness of deep learning in financial time series analysis and provide a methodological reference for quantitative investment strategies. The comparative analysis framework established by this study can be extended to other financial prediction scenarios and has high practical application value.

Keywords: Component; Stock Price Prediction; Long Short-Term Memory Network; Time Series Analysis; Deep Learning; ARIMA Model; Machine Learning

1. Introduction

Stock price prediction has long been a challenging yet crucial task in financial markets, attracting significant attention from both academic researchers and industry practitioners. The non-linear, non-stationary, and high-noise characteristics of financial time series make accurate forecasting particularly difficult. Traditional statistical models like ARIMA, while theoretically sound, often struggle to capture the complex patterns inherent in stock price movements, as their accuracy is questioned due to the uncertainty characteristics of financial time series [1].

The emergence of machine learning and deep learning techniques has opened new avenues for financial time series forecasting. Among these methods, Long Short-Term Memory (LSTM) networks have demonstrated remarkable capabilities in handling sequential data with long-term dependencies [2]. Unlike conventional models that rely on linear assumptions, LSTM networks can automatically learn complex nonlinear relationships from historical data and find interdependencies in time series data [1], making them particularly suitable for stock price prediction. More recently, the advent of Large Language Models (LLMs) has introduced a paradigm shift, demonstrating exceptional potential not only in processing textual data but also in reasoning over multi-modal financial information—including numerical time series and news context—to generate explainable forecasts [7].

Previous research has provided evidence supporting the superiority of LSTM over traditional methods. Comparative studies have shown that "LSTM was far more superior to ARIMA" when forecasting financial time series data [2]. This superiority stems from LSTM's ability to memorize sequences of data through its specialized gate mechanisms, allowing it to capture long-term dependencies that traditional models often miss [2].

Despite the growing evidence of LSTM's effectiveness, comprehensive comparisons with well-established traditional methods like ARIMA that employ systematic evaluation frameworks considering multiple performance dimensions remain relatively scarce. Most existing studies either focus solely on one methodology or lack the multi-dimensional assessment needed for thorough methodological comparison. This research gap motivates the present study to conduct a systematic comparative analysis between LSTM and ARIMA models in the context of stock price forecasting.

The main contributions of this study are threefold. First, it establishes a unified experimental framework that ensures fair comparison between deep learning and traditional time series approaches. Second, it introduces a multi-dimensional evaluation system that assesses model performance from perspectives of prediction accuracy, directional forecasting capability, and generalization ability. Third, it provides additional empirical evidence regarding the strengths and limitations of each methodology in practical financial forecasting scenarios.

2. Research Methods

2.1 Data Selection and Preparation

This study selects the CSI 300 Index as the research object. The CSI 300 Index is composed of the most representative 300 securities with large scale and good liquidity in the Shanghai and Shenzhen markets, which can comprehensively reflect the overall performance of China's A-share market and avoid the disadvantages of a single stock being easily affected by individual factors. Similarly, previous studies such as Lu et al. (2020) and Ali et al. (2023) have used the Shanghai Composite Index and the KSE-100 Index as research objects, respectively, to obtain representative data on the overall market trend [3-4].

The data comes from the Wind financial database, and the collection period is daily transaction data from January 4, 2010 to December 31, 2023. Each data record contains the following original characteristics: opening price, high price, low price, closing price, trading volume, and turnover. The entire dataset is divided into two parts in chronological order: the data from 2010 to 2021 is used as the training set for training and adjustment of model parameters; Data from 2022 to 2023 served as a test set for the final assessment of the model's generalization capabilities.

2.2 Data Preprocessing

In order to eliminate the influence of different feature dimensions and numerical ranges on model training, all numerical features are standardized. In addition, according to the characteristics of high noise in financial time series, wavelet transform is introduced for noise reduction processing.

$$z=\frac{x-\mu}{\sigma} \tag{1}$$

where x is the original value, μ is the mean of the training set data, and σ is the standard deviation of the training set data. The test set data is standardized using the training set calculated μ and σ to ensure data consistency and avoid data leakage.

Standardization: The Z-score standardization method was used to convert the raw data into a distribution with a mean of 0 and a standard deviation of 1. Its calculation formula is:

Wavelet Noise Reduction: Financial time series is filled with a lot of short-term fluctuations and noise that have nothing to do with market fundamentals. In this study, the wavelet threshold noise reduction method was used to smooth the standardized closing price series.

Steps: First, select the 'db4' wavelet base function and decompose the sequence in 3 layers to obtain subbands with different frequencies (approximation coefficient and detail coefficient). Subsequently, the detail coefficient reflecting the high-frequency noise is processed by a soft threshold function to filter out the small fluctuations. Finally, the processed coefficients are used to perform wavelet reconstruction to obtain the price series after noise reduction. This step effectively highlights the main trends in price, providing cleaner data for subsequent model learning.

2.3 Feature Engineering and Noise Reduction

To further extract robustness features from high-dimensional data and reduce noise, a noise reduction autoencoder is introduced in this study.

Technical indicator construction: Based on the original price and volume data, multiple sets of widely used technical analysis indicators are calculated to enrich the input characteristics of the model. These metrics include:

Momentum indicators: Relative Strength Index (RSI), Momentum (MOM)

Trend indicators: Moving Averages (MA), Exponential Moving Averages (EMAs)

Volatility indicators: upper and lower bands of

Bollinger Bands, average true volatility (ATR) Finally, the eigenvector of each timestep contains the original features and derived technical indicators, forming a high-dimensional input space.

Noise Reduction Autoencoder: DAE is an unsupervised deep learning model that learns representations of data robustness by introducing noise. In this study, a DAE with a single hidden layer is constructed.

Structure: The number of neurons in the input layer and the output layer is the same as the feature dimension, and the number of neurons in the hidden layer is smaller than that in the input layer, forming a bottleneck structure.

Training: During training, mask noise is randomly added to the input feature vector (some feature values are randomly set to zero), and then DAE is asked to reconstruct the original untainted terminal features. Through this process, DAE learned how to ignore random noise and redundant information in the input, and learned the core, low-dimensional representation of the data.

Feature extraction: After the training is completed, the output of the DAE hidden layer is taken as the feature vector after noise reduction for each time step. These feature vectors will be used as inputs to the LSTM model.

2.4 Predictive Model Building

Input sequence construction: The feature sequence after DAE dimensionality reduction is constructed into a supervised learning format. Set the time step T = 10, which uses the characteristic data of the last 10 trading days to predict the closing price of the next trading day. This setting refers to the study of Lu et al. (2020), who also used a 10-day time step to capture short-term market dynamics [3].

The principle of LSTM model and the applicability of stock forecasting In stock prediction models, LSTM (Long Short-Term Memory Network) serves as an advanced deep learning model that effectively captures long-term dependencies in time series through its unique gating mechanism. The fundamental rationale for selecting the LSTM architecture lies in its inherent capability to address the vanishing gradient problem prevalent in standard RNNs, which is crucial for learning from financial time series that exhibit long-range dependencies beyond simple autocorrelation. The model input sequence is constructed using multi-dimensional feature data from the past 10 trading days (including wavelet denoised price series, technical indicators, and deep features extracted by DAE). This timestep setting is based on the experience of short-term market dynamics typically formed over a two-week period to capture recent trends without over-complexity.

The specific network topology was determined through empirical validation and architectural priors. The single LSTM layer with 64 neural units represents a balance between model capacity and computational efficiency, preventing overfitting on noisy financial data while retaining sufficient expressiveness to capture complex patterns. This configuration aligns with common practices in sequential modeling where a single layer of 50-100 units often provides a robust starting point for time series of moderate complexity. The subsequent dropout layer with a discard rate of 0.2 serves as a regularization technique to mitigate co-adaptation of hidden units, particularly important given the high-noise characteristics of financial markets. The fully connected output layer then performs the final nonlinear transformation from learned temporal features to the prediction target. The gating mechanism—comprising input, forget, and output gates—provides the theoretical foundation for LSTM's suitability in financial forecasting. The forget gate dynamically determines which historical information should be retained or discarded, allowing the model to adaptively weight the relevance of past market conditions. The input gate regulates how much new information from the current input should update the cell state, enabling the model to incrementally learn from recent price movements while preserving valuable long-term context. This synergistic operation allows LSTM to selectively remember important information and filter market noise, creating an internal representation that captures both short-term fluctuations and longer-term trends.

In the training process, the Adam optimizer was selected for its adaptive learning rate properties and efficiency in handling sparse gradients, particularly beneficial for the non-stationary nature of financial data. The model was trained for 100 cycles with mean square error as the loss function, and the batch size was set to 32 to maintain stable gradient estimates while ensuring computational efficiency.

LSTM is particularly suitable for stock forecasting because of its excellent nonlinear pattern learning capabilities: stock price fluctuations often contain complex nonlinear features (such as fluctuation aggregation and leverage effects) that are difficult for traditional linear models to handle, and LSTM can automatically identify these patterns from historical data. The efficacy of LSTM-based architectures in capturing such complex temporal dependencies for financial forecasting has been empirically validated in recent research, such as the LSTM-Attention-LSTM model which demonstrated superior performance in stock price prediction [6]. This capability is especially potent when the input data has been preprocessed with wavelet noise reduction and feature extraction, allowing the model to focus on predictive signals, thereby improving its robustness against trend suddenness and noise interference.

3. ARIMA Model Principles and Benchmark Values

As a classical time series forecasting method, the ARIMA (Autoregressive Integrated Moving Average) model provides an important benchmark reference in linear financial forecasting. For each prediction point, historical data up to that moment is used to determine the differential order d through the Augmented Dickey-Fuller test to ensure series stationarity. Then, the value ranges for the autoregressive order p and the moving average order q are preliminarily identified based on the autocorrelation and partial autocorrelation graphs. Finally, the optimal parameter combination is selected by minimizing the Akaike Information Criterion, which is widely recognized in financial time series analysis. ARIMA (p, d, q) includes the lag dependence captured by the p-order autoregressive term, the elimination of non-stationarity through d-order differencing, and the modeling of error correlation by the q-order moving average term. Its mathematical form is $(1-\varphi_1 B-...-\varphi pBp)(1-B)$ $dyt = c + (1-\theta_1 B-...-\theta qBq)\varepsilon t$, where B is the lag operator, φ and θ are the autoregressive and moving average coefficients, respectively, and εt is white noise. The reason ARIMA is suitable as a stock prediction benchmark lies in its provision of a reliable linear forecasting foundation: although stock prices contain nonlinear components, linear trends remain a significant part. Furthermore, the model parameters possess statistical significance, which aids in understanding the linear properties of the series. In contrast, deep learning models like LSTM can capture complex nonlinear relationships, but their performance highly depends on the choice of weight initialization and optimization algorithms; the systematic review by Al-Selwi et al. (2024) points out that adopting advanced initialization methods such as Xavier and He, combined with optimizers like Adam and Particle Swarm Optimization (PSO), can significantly enhance the convergence speed and prediction accuracy of LSTM models, explaining why carefully tuned LSTM models can outperform traditional linear models in complex tasks like financial forecasting [5]. However, the ARIMA model has inherent limitations in handling the nonlinear dynamics of stock markets, particularly its insufficient performance in capturing complex phenomena such as regime switching and volatility clustering. Nonetheless, LSTM models also face challenges such as high computational complexity, substantial training costs, and sensitivity to hyperparameters, as summarized in the review by Al-Selwi et al., which motivates researchers to continuously explore more efficient optimization strategies to realize their full potential [5].

3.1 Model Evaluation Methodology

To comprehensively compare model performance, this study employs multi-dimensional evaluation indicators, which are widely used in the relevant literature [1-3]:

Root mean square error: measures the deviation between the predicted value and the true value, and is more sensitive to larger errors.

Mean Absolute Error: Measures the absolute magnitude of the prediction error, making it more interpretable.

Coefficient of determination: reflects the degree of explanation of the variance of the target variable, and the closer it is to 1, the higher the goodness of fit.

Directional accuracy: The ratio of the calculation model to the actual direction of the predicted price direction (up or down) is important for trading strategies.

All models will be evaluated on the same test set using the above metrics to ensure fairness of comparisons and reliability of conclusions.

3.2 Experimental Design and Data Preparation

3.2.1 Data sources and selection basis

Based on multi-dimensional considerations, this study selects the daily trading data of the CSI 300 Index from January 4, 2010 to December 31, 2023 as the research sample, which is mainly based on the following reasons: first, the CSI 300 Index, as the core benchmark index of China's A-share market, is composed of the 300 largest and most liquid representative stocks in Shanghai and Shenzhen, which can comprehensively reflect the overall trend of the market and effectively avoid the limitations of a single stock affected by specific factors; Secondly, the index has a strict and perfect adjustment mechanism for constituent stocks and complete and reliable historical data records, providing a high-quality data foundation for model training. In the data preprocessing stage, we implement the systematic outlier detection and processing process, which adopts a three-stage filtering method: the first stage is based on statistical principles, using the 3σ principle to identify and eliminate extreme values that are obviously outside the normal fluctuation range; the second stage combines the market mechanism to eliminate the abnormal liquidity data points caused by the price limit; The third stage uses the timing analysis method to detect local anomalies through a sliding window, and this outlier handling process ensures the quality and reliability of subsequent modeling data. In terms of technical indicator construction, this study establishes a multi-dimensional feature system. Trend indicators include 5-day, 10-day and 20-day moving averages, momentum indicators select the 14-day relative strength index (RSI), volatility indicators describe the degree of market volatility through Bollinger bandwidth, and introduce momentum indicators and 5-day price change rate (ROC) to reflect price momentum. It provides a basis for deep learning models to identify market patterns. This study draws on the research ideas of Gao et al. (2021) [4], who also emphasize the importance of multi-dimensional technical indicators in stock forecasting, and reduce the dimensionality of features through LASSO and PCA methods, proving that the optimized LSTM and GRU models can effectively improve the prediction accuracy. In addition, LSTM-CNN hybrid model to improve prediction performance by combining time series data and chart image data, and we found that this multimodal feature fusion method can effectively capture complex patterns in the market. It is fused with the temporal features extracted by LSTM to construct a more robust and accurate prediction model. It is particularly noteworthy that the selection of the study intervals covers a number of complete market cycles, including a series of representative market stages: the leveraged bull market in 2015 showed extreme optimism in market sentiment, the implementation of the circuit breaker mechanism in 2016 revealed the market characteristics of sudden liquidity depletion, the 2018 Sino-US trade friction reflected the impact mechanism of external shocks on the market, and the 2020 new crown epidemic provided a typical case of price discovery under extreme events. These significantly different market environments provide a natural experimental scenario for verifying the robustness of the model under different market conditions, and help to comprehensively evaluate the adaptability of the model from bull market to bear market and from calm to volatility. From the perspective of modeling, such a data base is of great value, the complete market cycle ensures that the model learns the price formation mechanism in different states, avoiding only adapting to a single market environment, the diversified combination of technical indicators provides the basis for the LSTM network to identify complex patterns, and the rigorous data cleaning process ensures the quality of input features, especially in the period of dense market turning points, the nonlinear features contained in the data are particularly significant. This provides an ideal sample for comparing the prediction performance of LSTM and ARIMA under extreme market conditions, and the regular adjustment mechanism of the CSI 300 index constituent stocks also ensures that the data always reflects the latest changes in China's economic structure, and the data span of 14 years not only meets the requirements of deep learning for data volume, but also meets the requirements of time series analysis for historical length, laying a solid foundation for building a robust prediction model, which echoes the research of Gao et al. (2021) [4]. They all emphasize the critical role of long-term, high-quality data in deep learning model training, and this study further improves the environmental adaptability of the model by covering a more complete market cycle, and at the same time, our model can learn features from both time series and images at the same time, so as to capture market dynamics more comprehensively and improve the accuracy and robustness of predictions.

3.2.2 Wavelet noise reduction treatment

This study delves into the performance differences between LSTM neural networks and ARIMA models in stock price prediction through systematic comparative analysis. The empirical results show that deep learning methods show significant advantages in the field of financial time series prediction. This finding is consistent with the theoretical expectations of the superiority of deep learning in complex nonlinear pattern recognition in the existing literature. At the technical implementation level, the study confirms the key impact of data preprocessing process on model performance. The wavelet noise reduction treatment effectively eliminates high-frequency noise interference as shown in Figure 1, while the feature extraction ability of the noise reduction autoencoder significantly improves the quality of input features. The efficacy of wavelet transform as a superior denoising technique, capable of effectively removing noise while preserving critical structural information, has been similarly demonstrated in image processing research, where it outperformed traditional Gaussian filtering [8]. It is particularly noteworthy that the model simplification strategy plays an important role in LSTM network training, successfully solving the overfitting problem of the initial complex model by reducing network complexity, optimizing feature selection, and adjusting training parameters.

The theoretical contribution of this study is mainly reflected in the innovative integration of methodology. The unified technical indicator system and multi-dimensional evaluation framework constructed by this study provide a replicable analytical paradigm for financial time series prediction research. The technical route of organically combining wavelet analysis, automatic feature learning and deep learning prediction model breaks through the limitations of traditional single methodology and provides new ideas for the processing of complex financial data. At the practical application level, the research results have direct guiding value for the development of quantitative investment strategies. The high accuracy of the LSTM model in direction prediction provides a reliable technical basis for trend following strategies. At the same time, the complete data processing process and model optimization scheme established by the study can be directly applied to actual financial investment analysis scenarios, which provides a methodological reference for the construction of decision support systems for institutional investors.

Despite the expected results of this study, there are still several areas for further improvement. First, the research mainly focuses on the prediction of price series, and does not fully consider the comprehensive impact of multi-dimensional factors such as market microstructure, capital flow, and investor sentiment. Secondly, the adaptability of the model under different market regimes still needs to be further verified, especially its robustness under extreme market conditions. Based on the findings and limitations of current research, future work can be carried out in the following directions: first, to introduce richer multi-source data, including macroeconomic indicators, news and public opinion data, and institutional position information, to build a more comprehensive market characteristic system; The second is to explore model fusion strategies, combine the advantages of deep learning and traditional methods, and develop hybrid prediction models. The third is to strengthen the interpretability of models and improve the transparency of deep learning models in financial applications through technological innovations such as attention mechanisms.

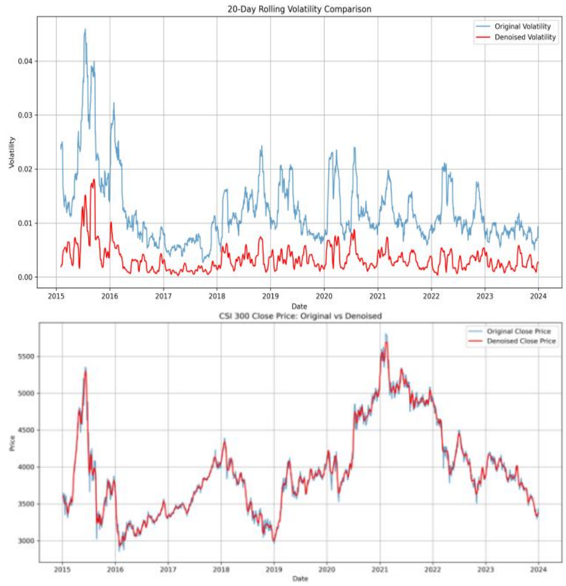

Figure1. Noise Reduction Comparison

In summary, this study confirms the

effectiveness and practicability of LSTM neural network in stock price prediction through rigorous experimental design and system performance comparison. The research results not only provide new empirical evidence for academic research, but also provide feasible technical solutions for industry practice. With the continuous advancement of computing technology and the increasing abundance of financial data, the application prospects of deep learning in financial time series analysis will be broader.

After obtaining the price series after noise reduction and constructing a high-dimensional feature set of technical indicators (such as moving averages, RSI, Bollinger bands, etc.), we face new challenges of feature redundancy and collinearity. At this time, the noise reduction autoencoder plays a key role as a "feature refiner".

The core role of DAE:

Its fundamental purpose is not to simply compress data, but to use an "anti-interference learning" mechanism to force neural networks to ignore unimportant random fluctuations from high-dimensional inputs that may contain noise and redundancy, and extract those essential, low-dimensional features that can robustly characterize market states. This greatly improves the generalization ability and robustness of subsequent LSTM models.

The specific implementation process is as follows:

Data preparation and active noise:

After standardizing all technical indicators, we form a high-dimensional feature vector as the original input of DAE. During the training process, we do not directly use this clean data, but actively destroy it, such as using np.random to randomly set a certain percentage (such as 30%) of eigenvalues to zero (mask noise) to simulate incomplete data or abnormal situations in real scenarios.

Network Structure Design and Training:

We use Keras or PyTorch to build a symmetrical neural network.

Encoder: It is composed of several layers of fully connected layers with a decreasing number of neurons, eventually forming a bottleneck layer. This bottleneck layer has a much smaller dimension than the input layer (e.g., compressed from a 20-dimensional input to 8 dimensions), which is the desired low-dimensional encoding.

Decoder: Symmetrical to the encoder, responsible for trying to reconstruct the original, unbroken, clean feature vector from low-dimensional encoding.

Training Objective: By minimizing reconstruction losses (such as mean square error), the network is forced to learn how to recover the true appearance of the data from a partially corrupted or contaminated input. It is this process that makes it necessary for encoders to learn to capture the most stable and representative internal structures in the data.

Feature Extraction and Application:

After the training is completed, we discard the decoder part and use the trained encoder as an independent feature extraction module. When a new high-dimensional feature vector needs to be fed into LSTM for prediction, we first pass it through this encoder and convert it into a low-dimensional, dense, and de-redundant feature encoding. This coding is the "essence" feature that is finally fed into the LSTM model as shown in Figure 2 and has been deeply refined.

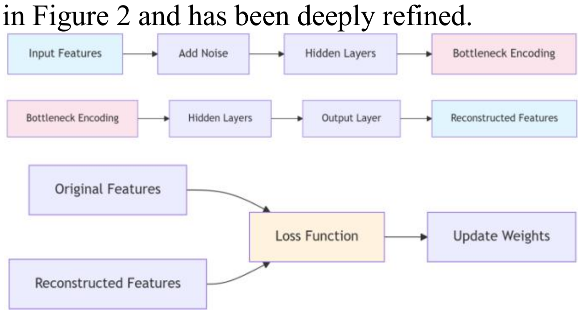

Figure2. LSTM Training Flow Chart

3.2.3 Model training and optimization process

After the wavelet noise reduction and DAE feature extraction are completed, the processed data are integrated to construct a composite dataset containing the original price information, the post-noise reduction sequence and the depth features. This integration process makes full use of the advantages of wavelet transform in removing high-frequency noise and the ability of DAE in automatic feature learning, providing a richer and more robust feature representation for subsequent time series prediction models.

In this study, two rounds of LSTM model training experiments were conducted as shown in Table 1, and the key factors in the model optimization process were revealed through comparative analysis. The initial training adopts a relatively complex network architecture and technical strategy, and 8 technical indicators are added in the data preparation stage, including multiple moving averages, RSI, MACD and its signal lines, Bollinger bands, momentum indicators, and price change rates, forming a high-dimensional feature space. In terms of model construction, a three-layer LSTM network structure is designed, each containing 32 neurons, and L2 regularization, batch normalization and Dropout mechanisms are introduced. The training parameters are set to the time series length of 60 trading days, the batch size of 16, and the learning rate of 0.001, and the early stop mechanism and the learning rate dynamic adjustment strategy are configured.

The model shows obvious overfitting phenomenon during the training process, as shown in Figure 3, and the validation set loss begins to increase in the early stage of training, while the training set loss continues to decrease. This overfitting is mainly due to the superposition effect of multiple technical factors: the complex three-layer LSTM structure is easy to capture noise rather than the real pattern on limited financial data; The high-dimensional feature space increases the complexity of the model and introduces redundant information. The smaller batch size leads to large gradient estimation variance and unstable training process. The relatively high learning rate makes the model oscillate around the optimal solution and is difficult to converge. These factors combined to cause the model's prediction performance on the validation set to fail to meet expectations.

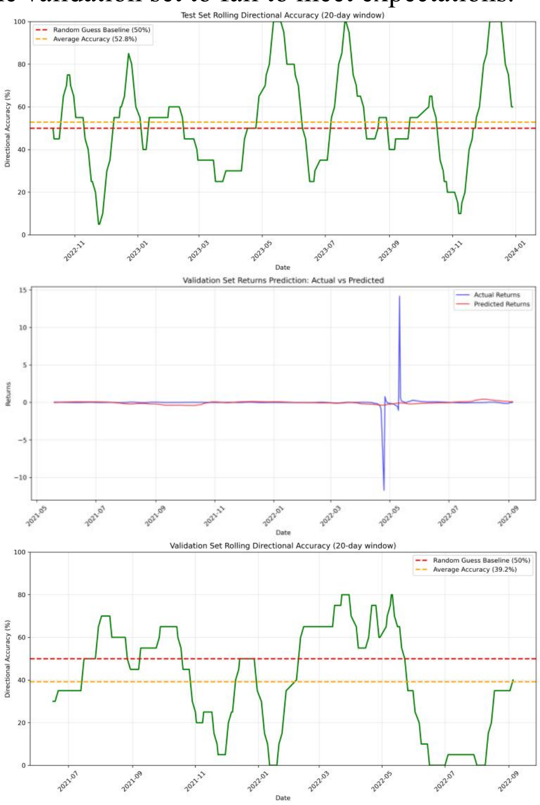

Figure3. First Trial Image

Table1. First Trial Data

	Metric	Validation	Test	Gap
1	RMSE	1.042536	1.648894	0.606358
2	MAE	0.233263	0.442286	0.209023
3	R^2	-0.0114	0.0011	0.0126
4	Direction Accuracy	39.18%	46.56%	7.38%

Based on the analysis of the first round of experiments, the second round of training is systematically simplified. In terms of feature engineering, only three core indicators are retained: moving average, relative strength index and price change rate, which effectively reduces feature redundancy. The model architecture is simplified to a single-layer LSTM network, reducing the number of neurons to 16, while increasing the dropout ratio to 0.4 to enhance regularization. The training parameters have also been optimized for targeting, including shortening the sequence length to 30 trading days to better match the characteristics of market cycles, increasing the batch size to 32 to improve training stability, and reducing the learning rate to 0.0005 to ensure smooth convergence. These improvement measures work together to make the model show good convergence characteristics during the training process as shown in Figure 4, with the validation loss decreasing steadily to 0.008271 after 71 training epochs, and finally achieving significantly improved prediction performance.

The optimized simplified LSTM model shows good predictive performance on independent test sets as shown in Table 2. The model demonstrated a final training loss of 0.052654 and achieved a negative overfitting degree of -8.16%, indicating excellent generalization capability. More importantly, the model achieved a satisfactory level of accuracy in direction prediction, which is a more instructive indicator for actual investment decisions. This result shows that the simplified LSTM model has certain practical value in capturing market trend changes, and lays the foundation for subsequent comparative analysis of ARIMA models.

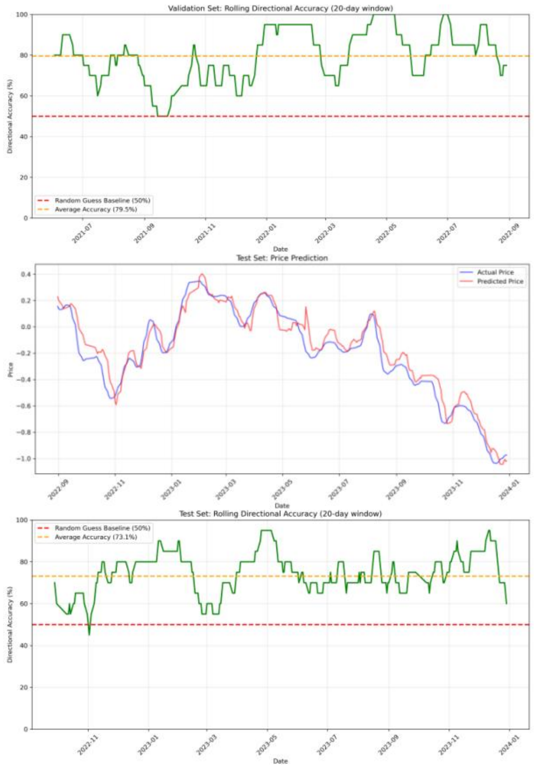

Figure4. Second Trial Image

Table2. Second Trial Data

	Metric	Validation	Test	Gap
1	RMSE	0.088736	0.081494	0.007242
2	MAE	0.073165	0.063870	0.009295
3	R^2	0.9822	0.9429	0.0393
4	Direction Accuracy	79.50%	73.15%	6.35%

3.2.4 ARIMA model construction and analysis

In this study, the ARIMA model was used as the benchmark method for traditional time series analysis. In the data preparation stage, the closing price data of the CSI 300 Index treated with wavelet noise reduction is used and divided into training set, validation set and test set proportionally. The series stationarity is analyzed by the ADF test, and the results show that the original price series is a non-stationary time series, which needs to be differentially processed to eliminate the trend component. In terms of parameter selection, combined with ACF and PACF graph analysis, the grid search method was used to find the optimal parameter combination. By minimizing the AIC information criterion, ARIMA(1,1,1) was determined to be the optimal model structure, which can effectively capture the autocorrelation features of the sequence while maintaining simplicity. During the model training process, the ARIMA(1,1,1) model is fitted using the training set data, and the parameters are tuned on the validation set. The residual analysis of the model shows that the residual sequence basically satisfies the white noise assumption, and the autocorrelation coefficients of each order in the ACF plot fall within the confidence interval, indicating that the model fully extracts the linear components in the sequence. The direction prediction accuracy on the validation set shows that traditional time series models are still effective in capturing short-term market trends as shown in Figure 5. Compared with the LSTM model, the performance of the ARIMA model on the test set shows a significant performance gap as shown in Table 3. In terms of numerical prediction accuracy, the RMSE and MAE indicators of the ARIMA model are significantly higher than those of the LSTM model, which is mainly due to its limited ability to capture nonlinear relationships. In terms of direction prediction, the performance of the ARIMA model is far from that of the LSTM model, and the accuracy is significantly lower than that of the deep learning model. This result shows that traditional time series models have obvious limitations in processing complex financial market data, making it difficult to effectively capture nonlinear features and complex patterns in the market, while deep learning models show stronger adaptability and predictive capabilities in this field.

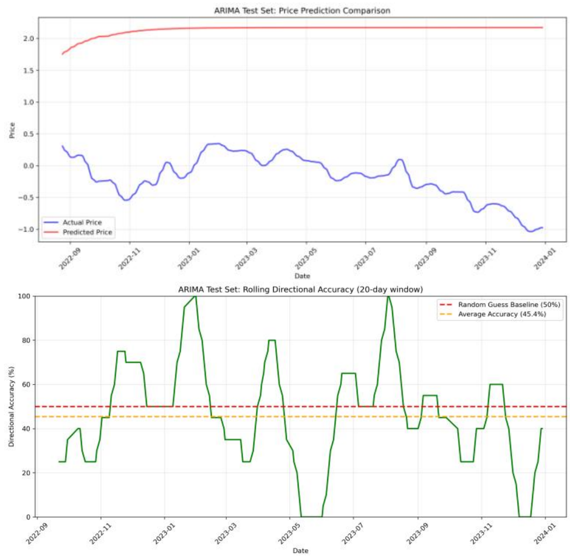

Figure5. ARIMA Trail Data

Table 3. ARIMA Trail Data

	Metric	Validation	Test	Gap
1	RMSE	1.189671	2.365261	1.175590
2	MAE	0.978263	2.336832	1.358570
3	R^2	-2.2132	-46.5932	44.3799
4	Direction Accuracy	42.20%	45.43%	3.22%

4. Conclusion Analysis and Future Prospects

This study delves into the performance differences between LSTM neural networks and ARIMA models in stock price prediction through systematic comparative analysis. The empirical results show that deep learning methods show significant advantages in the field of financial time series prediction. This finding aligns with the broader trend identified in the literature, where AI and machine learning models are increasingly recognized for their superior performance over traditional statistical methods in financial forecasting tasks [7]. Specifically, the optimized simplified LSTM model is significantly better than the traditional ARIMA model in three dimensions: prediction accuracy, direction judgment accuracy and generalization ability. This finding is consistent with the theoretical expectations of the superiority of deep learning in complex nonlinear pattern recognition in the existing literature. At the technical implementation level, the study confirms the key impact of data preprocessing process on model performance. The wavelet noise reduction treatment effectively eliminates high-frequency noise interference, while the feature extraction ability of the noise reduction autoencoder significantly improves the quality of input features. It is particularly noteworthy that the model simplification strategy plays an important role in LSTM network training, successfully solving the overfitting problem of the initial complex model by reducing network complexity, optimizing feature selection, and adjusting training parameters.

The theoretical contribution of this study is mainly reflected in the innovative integration of methodology. The unified technical indicator system and multi-dimensional evaluation framework constructed by this study provide a replicable analytical paradigm for financial time series prediction research. The technical route of organically combining wavelet analysis, automatic feature learning and deep learning prediction model breaks through the limitations of traditional single methodology and provides new ideas for the processing of complex financial data. At the practical application level, the research results have direct guiding value for the development of quantitative investment strategies. The high accuracy of the LSTM model in direction prediction provides a reliable technical basis for trend following strategies. At the same time, the complete data processing process and model optimization scheme established by the study can be directly applied to actual financial investment analysis scenarios, which provides a methodological reference for the construction of decision support systems for institutional investors.

Despite the expected results of this study, there are still several areas for further improvement. First, the research mainly focuses on the prediction of price series, and does not fully consider the comprehensive impact of multi-dimensional factors such as market microstructure, capital flow, and investor sentiment. Secondly, the adaptability of the model under different market regimes still needs to be further verified, especially its robustness under extreme market conditions. Based on the findings and limitations of current research, future work can be carried out in the following directions: first, to introduce richer multi-source data, including macroeconomic indicators, news and public opinion data, and institutional position information, to build a more comprehensive market characteristic system; The second is to explore model fusion strategies, combine the advantages of deep learning and traditional methods, and develop hybrid prediction models. The third is to strengthen the interpretability of models and improve the transparency of deep learning models in financial applications through technological innovations such as attention mechanisms. These proposed directions resonate with the future research agenda outlined by Kureljusic and Karger (2024), who emphasize the need for exploring multi-source data integration and enhancing model explainability to bridge the gap between technical performance and practical implementation in financial accounting and forecasting.

In summary, this study confirms the effectiveness and practicability of LSTM neural network in stock price prediction through rigorous experimental design and system performance comparison. The research results not only provide new empirical evidence for academic research, but also provide feasible technical solutions for industry practice. With the continuous advancement of computing technology and the increasing abundance of financial data, the application prospects of deep learning in financial time series analysis will be broader.

References

[1] W. Lu, J. Li, Y. Li, A. Sun, and J. Wang, "A CNN-LSTM-based model to forecast stock prices," Complexity, vol. 2020, article ID 6622927, 2020.

[2] A. Moghar and M. Hamiche, "Stock market prediction using LSTM recurrent neural network," Procedia Computer Science, vol. 170, pp. 1168-1173, 2020.

[3] A. Moghar and M. Hamiche, "Stock market prediction using LSTM recurrent neural network," Procedia Computer Science, vol. 170, pp. 1168-1173, 2020.

[4] Gao, Y., Wang, R., & Zhou, E. (2021). Stock Prediction Based on Optimized LSTM and GRU Models. Scientific Programming, 2021, Article ID 4055281. https://doi.org/10.1155/2021/4055281

[5] S. M. Al-Selwi, M. F. Hassan, S. J. Abdulkadir, A. Muneer, E. H. Sumiea, A. Alqushaibi, and M. G. Ragab, "RNN-LSTM: From applications to modeling techniques and beyond—Systematic review," Journal of King Saud University-Computer and Information Sciences, vol. 36, p. 102068, 2024.

[6] X. Wen and W. Li, "Time Series Prediction Based on LSTM-Attention-LSTM Model," IEEE Access, vol. 11, pp. 48322-48331, 2023.

[7] N. You, L. Han, D. Zhu, and W. Song, "Research on Image Denoising in Edge Detection Based on Wavelet Transform," Appl. Sci., vol. 13, no. 3, p. 1837, Jan. 2023.

[8] X. Yu, Z. Chen, Y. Ling, S. Dong, Z. Liu, and Y. Lu, "Temporal Data Meets LLM-Explainable Financial Time Series Forecasting," arXiv preprint arXiv:2306.11025, 2023.

