Applied Research on Gold Stock Price Prediction Based on Long Short-Term Memory Neural Network

Jikai Zhang [1,*], Suiyu Yang [2]

[1] Department of Digital Economics, Shanghai University of Finance and Economics, Shanghai, China, 200433

[2] School of Information Science and Technology (School of Cyber Security), Guangdong University of Foreign Studies, Guangzhou, China, 510006

* Corresponding Author Email: zjk87068353@163.com

Abstract. As financial markets become increasingly intricate, accurately predicting gold stock prices has become an indispensable decision - making tool for investors. This not only aids in risk management but also holds significant potential for profit - seeking in the volatile gold stock market. This paper proposes a gold stock price prediction model based on LSTM. By integrating multidimensional technical indicators such as Moving Average (MA) and Relative Strength Index (RSI), it optimizes feature engineering to effectively capture price fluctuation patterns. To tackle the hyperparameter sensitivity issue during LSTM training, an exhaustive search method is employed for hyperparameter optimization. This meticulous approach helps the model to better adapt to various market scenarios and historical data characteristics, thus enhancing its generalization ability. The optimized LSTM model shows remarkable improvements in performance metrics including Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE), validating its effectiveness in capturing both long - term trends and short - term volatility.

Keywords: Long Short-Term Memory Neural Network (LSTM), Gold Stock Price Prediction, Time Series Analysis, Model Optimization, Financial Forecasting.

1. Introduction

As global financial markets grow more dynamic and the economy evolves at a rapid pace, predicting stock market trends has emerged as a focal point for investors and financial analysts. In the context of heightened global economic uncertainty, gold, as a safe - haven asset, exerts an increasingly significant influence on stock markets through its price fluctuations.

Neural network models, functioning as nonlinear dynamic learning systems, can execute effective and high - precision predictive fitting. By learning from historical characteristic data of changes, they approximate evolutionary patterns, thus reflecting future trends to some extent. In the volatile stock market environment, neural network technology has been widely embraced for stock market prediction. This is attributed to its outstanding learning capabilities and robust simulation functions. Compared with classical econometric methods, neural networks demonstrate distinct advantages in financial time - series analysis.

In the realm of financial academic research, the application of neural network technology in stock market prediction mainly centers on nonlinear dynamic system modeling and time - series analysis. The stock market, a complex nonlinear dynamic system, has its nonlinear price fluctuation patterns and time - series characteristics as the core of research. Neural networks, with their nonlinear approximation capabilities and adaptive features, are extensively used to identify and predict the laws governing stock price movements. Among them, recurrent neural networks (RNNs) and their variants, such as LSTM and gated recurrent units (GRU), show remarkable prowess in processing time - series data. They achieve this by capturing long - term dependencies within sequences, thereby enhancing prediction accuracy.

This study employs a LSTM model. The model is capable of flexibly storing and discarding information across different time steps, making it well-suited for processing and predicting critical events characterized by significant time intervals and delays within time-series data. When handling long sequences, it effectively addresses the issues of vanishing and exploding gradients, thereby capturing the long-term dependencies inherent in the data.

By constructing an LSTM - based prediction model for gold stock prices, this paper aims to: first, assess the effectiveness of the LSTM model in financial market prediction; second, explore and evaluate the model's ability to predict gold stock prices; and finally, provide investors with a novel tool to enhance their understanding of and predictive ability for stock market changes. This research not only advances the application of LSTM models in financial forecasting but also aids investors and decision - makers in better navigating the complexities of financial markets.

2. Literature Review

Predicting stock prices is a complex task involving multiple methods and technologies, and scholars worldwide have extensively explored the application of neural networks in the stock market. While discussing algorithms for stock market prediction, apart from LSTM, widely adopted methods include RNN, ARIMA, SVM, and XGBoost. Traditional models like ARIMA adapt to time-series characteristics through differencing and autoregressive moving average processes; SVM relies on kernel functions for nonlinear mapping in regression tasks; XGBoost improves accuracy by constructing gradient-boosted decision trees. However, these methods face challenges in processing highly volatile nonlinear data, prone to overfitting and limited in capturing long-term dependencies.

As an advanced RNN variant, LSTM addresses the gradient vanishing issue via its gating mechanisms (input, forget, and output gates), enabling effective learning of long-term dependencies in time-series data. This advantage has made it prominent in financial prediction: Mo Jiawei [1] demonstrated its superior accuracy in predicting Shanghai 50 constituent stocks compared to traditional networks; Yang Qing et al. [2] showed that LSTM outperformed SVR, MLP, and ARIMA in both short- and long-term forecasts.

Domestic scholars have further improved LSTM through hybrid approaches. Cui Jianmin [3] simplified model complexity via a siamese-LSTM network; Dou Wei [4] enhanced prediction effectiveness with an Attention-LSTM model; Ren Jun et al. [5] improved generalization of LSTM via regularization for Dow Jones Index forecasting; Song Gang et al. [6] optimized LSTM with PSO, achieving higher accuracy and universality; Peng Yan et al. [7] boosted prediction accuracy by 30% using L2 regularization and dropout. Hybrid models like LSTM-CNN-CBAM [8] exhibit better timeliness, though Hu Yamei et al. [9] cautioned that excessive hybridization does not guarantee higher accuracy.

Notably, existing research primarily relies on technical indicators (e.g., price, volume) as inputs, neglecting critical external factors such as government policies and market sentiment. While studies like Liao Chang [10] and Yu Xiaojian et al. [11] have integrated textual sentiment features to enhance prediction accuracy, the systematic quantification of multi-dimensional influences remains underdeveloped. This gap highlights the need for more comprehensive feature engineering that bridges numerical data and qualitative market dynamics, providing a critical direction for improving model robustness in real-world financial scenarios.

3. Model building

3.1. Basic concepts of the model

LSTM, a specialized type of RNN, is inherently adept at processing and predicting time-series data. By introducing gating mechanisms, LSTM addresses the vanishing gradient problem that plagues traditional RNNs when handling long-sequence data, enabling it to effectively capture long-term dependencies. An LSTM unit receives the previous hidden state (ht−1), cell state (Ct−1), and current input (xt), where the forget gate (ft) determines old information to discard, and the input gate (it) governs new information to update. Through these gating operations, the cell state is updated, and the current hidden state (ht) is computed via the output gate (ot). In this study, the structural design of LSTM is leveraged to explicitly mitigate the vanishing gradient issue common in deep networks, thereby enhancing its generalization capability in gold stock price prediction, optimizing the training process, and more effectively modeling long-term dependencies within stock price data. Figure 1 illustrates the unit structure described above.

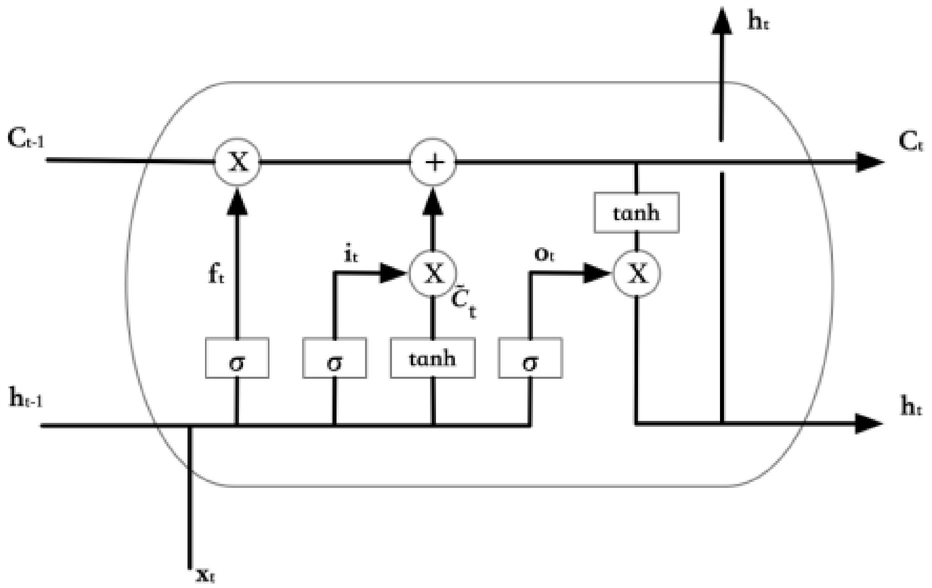

Figure 1. Schematic of LSTM Cell Operations

As depicted in Figure 2, gold stock price data are first obtained through an interface, followed by data cleaning and normalization of the closing price series. Subsequently, a time-series dataset is constructed using a predefined look_back parameter, reshaping the data into the three-dimensional array format required by the LSTM network. The model is then compiled using the Adam optimizer and MSE as the loss function. Normalized time-series data are employed to train the model for predicting future closing prices, with iterative parameter adjustment and optimization performed to generate the final predicted gold stock price data.

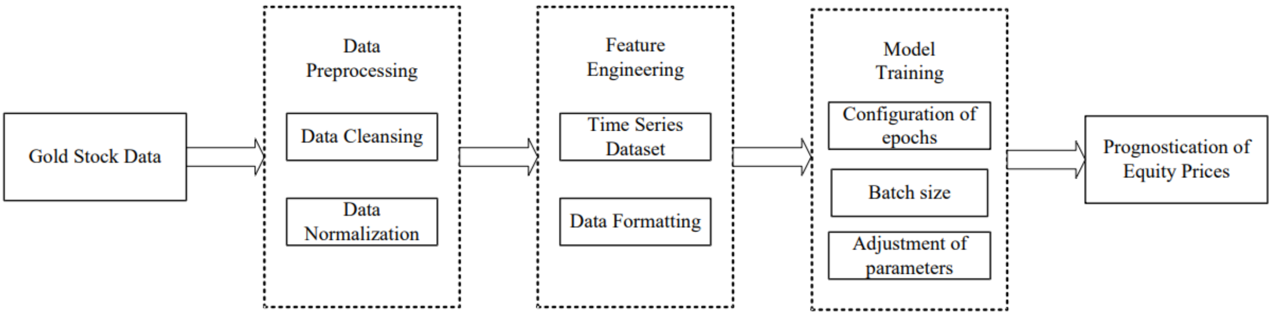

Figure 2. Workflow of LSTM Model for Gold Stock Price Prediction

3.2. Development of the LSTM Model

3.2.1. Algorithm Flow

The LSTM algorithm flow commences with data normalization, where closing price data are normalized to ensure values are scaled within the range of 0 to 1. In the feature engineering stage, multiple technical indicators—including moving average, moving average convergence divergence (MACD), and relative strength index (RSI)—are introduced as input features.

Prior to model construction, a time-series dataset is created through custom function implementation and parameter setting, followed by reshaping the dataset into a three-dimensional format compatible with the LSTM network. After completing these preparatory steps, the LSTM model is constructed with corresponding parameters and compiled. Subsequently, the model is trained after setting parameters for epochs, batch size, and verbose output.

The algorithm flow depicted in Figure 3 encapsulates the full lifecycle of a machine learning project—from data preprocessing to model training—with particular emphasis on data normalization, feature engineering construction, definition of LSTM and Dense layers, and parameter setting in model training.

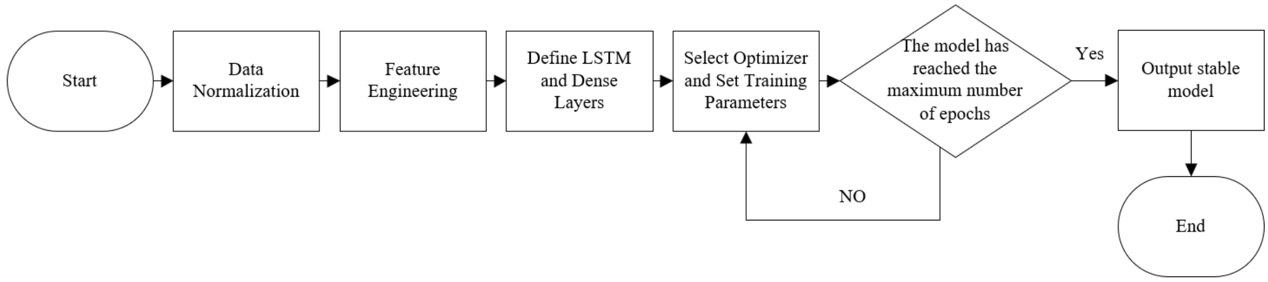

Figure 3. Schematic Diagram of the LSTM Algorithm Process

3.2.2. Model Training

In this section, RSI and MACD indicators are incorporated. As a momentum indicator measuring overbought or oversold stock price conditions, RSI, together with MACD—used to identify trend changes and momentum—collectively provide the model with quantified information on market sentiment and price momentum; their combined inclusion enhances the model's sensitivity to market dynamics. To comprehensively capture price volatility, Bollinger Bands are also constructed, offering additional information on price volatility to identify potential reversal points. Finally, all features are normalized to ensure uniform scaling, meeting the input requirements of the LSTM model and improving training stability and efficiency.

Raw data are first normalized to the 0–1 range, reducing the impact of dimensionality and enhancing training efficiency and effectiveness. Training data are then generated from the normalized dataset: specifically, data are partitioned along the time series, with each input sample containing five consecutive time steps of feature data, targeting the prediction of the next time step. This process establishes a mapping between inputs and outputs, laying the foundation for subsequent model training.

Subsequently, the generated training data are reshaped to fit the input requirements of the LSTM network, whose input layer expects three-dimensional data—with dimensions representing sample number, time steps, and feature number, respectively. The data are thus reshaped into the corresponding three-dimensional format to ensure proper input into the LSTM network.

During training, parameters for epochs (number of training iterations) and batch size (number of samples per training iteration) are set. Rational configuration of these parameters balances model learning effectiveness and training efficiency. Meanwhile, detailed output information from the training process is obtained via relevant parameters, enabling real-time monitoring of training status, timely identification of issues, and effective model adjustment—ensuring the model accurately learns data patterns and achieves valid predictions of target values.

4. Experiments

4.1. Data Preprocessing

This paper obtained data on Chinese gold stock prices from 2021 to 2024 through web crawlers, with key data features including date, opening price, closing price, highest price, lowest price, and trading volume. To enable the data to be more efficiently trained by the model, data preprocessing was performed. Specifically, after completing data cleaning, the data was sorted by date to ensure the continuity of the time series, and finally data normalization was conducted. Through the above processes, the preprocessing of the data was completed.

It was found that the "date" column in the data was not in a standard date format, which could inconvenience subsequent time - series analysis. Therefore, this paper decided to convert it to a unified date format—a logic based on considerations of data consistency and analytical requirements, providing convenience for subsequent operations such as time ordering and time - series charting. Further analysis revealed that the "trading volume" column was stored as a string, with some data containing unit identifiers such as "M" or "K". This non - standard storage format would interfere with numerical calculations and analysis. Driven by the pursuit of data accuracy and analytical precision, a method was designed to standardize the trading volume data, enabling it to participate in subsequent analysis as unified numerical values.

As shown in Figure 4, after completing data format conversion and standardization, the data was sorted by "date"—a logical step rooted in the inherent requirements of time - series data, where data points must be ordered chronologically to accurately reflect trends over time.

Figure 4. Candlestick chart

As shown in Figure 5, moving averages serve as critical input features in the LSTM model,. These features assist the model in learning the time - series characteristics of price data, thereby enabling more accurate predictions of future stock prices. When a short - term moving average crosses above a long - term moving average, it is typically regarded as a buy signal, which can be used as one of the bases for model training.

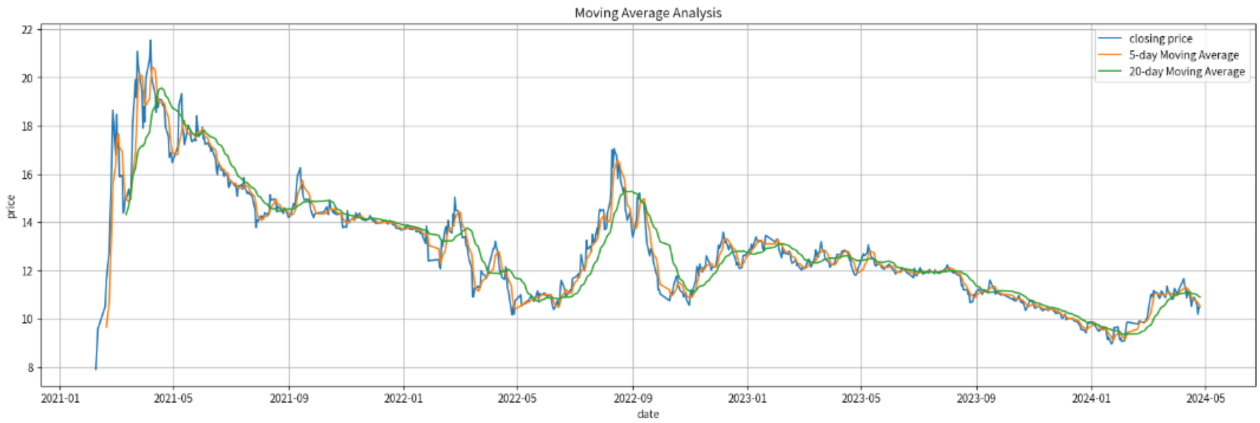

Figure 5. Moving Average Analysis

The RSI reflects market sentiment, providing sentiment-related features for the LSTM model to enhance its accuracy in predicting gold stock price trends. As illustrated in Figure 6, RSI values are calculated by computing differences in closing prices and moving averages of gains and losses, then plotted the RSI indicator chart. The chart displays the evolution of RSI values over dates, with two reference lines at 70 and 30 added to represent overbought and oversold conditions, respectively.

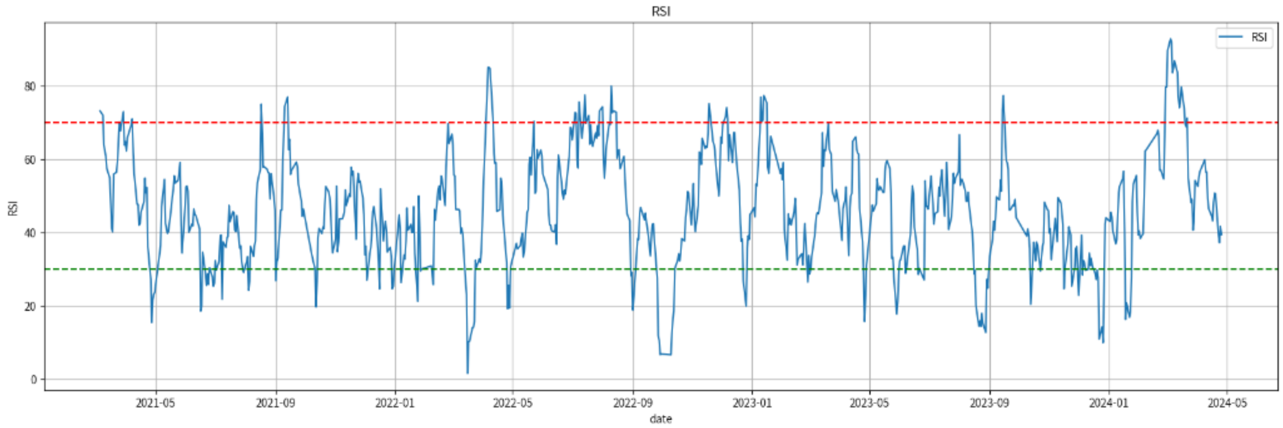

Figure 6. Relative Strength Index

Bollinger Bands, a volatility indicator based on moving averages and standard deviation, consist of three lines: the middle band is the 20-day moving average of closing prices, while the upper and lower bands are the middle band plus and minus twice the standard deviation, respectively. The width of Bollinger Bands changes with price volatility; when prices touch the upper or lower band, it often signals a potential entry into overbought or oversold states, providing important reversal signals for investors. In Figure 7, by calculating the three bands of Bollinger Bands, the chart visually depicts how stock prices move within the Bollinger Band channel, offering the LSTM model references for price volatility ranges and trend directions. The information on price fluctuation ranges and trend directions helps the model better grasp the stock price trajectory under different market conditions, especially providing critical reversal signals when prices reach key positions such as the upper or lower Bollinger Band, thus avoiding biases caused by the model over-relying on single price sequence information during prediction.

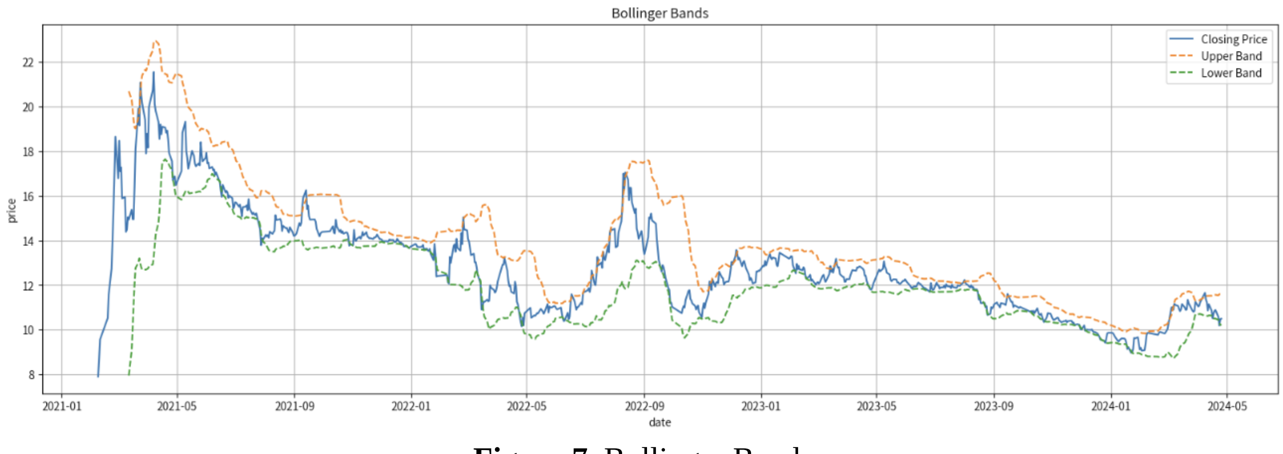

Figure 7. Bollinger Bands

The MACD indicator reflects medium-term trend changes in stock prices by calculating the difference between two exponential moving averages (EMA) of different periods. As shown in Figure 8, crossovers between the MACD line and the signal line, as well as changes in the length and color of the histogram, contain rich market information such as trend strength and momentum changes. These provide the LSTM model with dynamic features of medium-term price trends, enabling the model to better capture important market behaviors like medium-term trend reversals and accelerations, thereby improving prediction accuracy and stability.

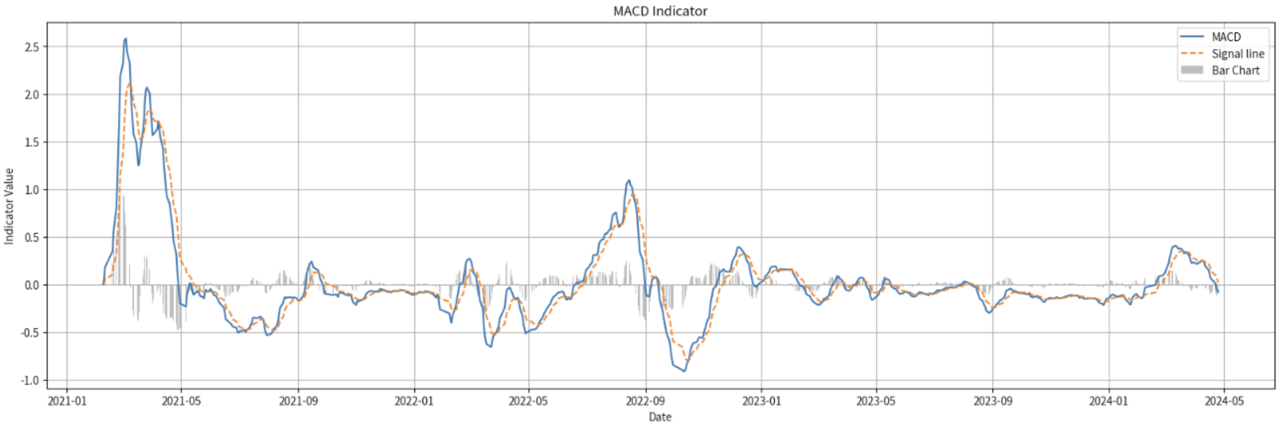

Figure 8. MACD Indicator

4.2. Evaluation Metrics

In this study, three key statistical indicators were employed during the model evaluation stage to quantify the performance of the LSTM prediction model: mean squared error (MSE), root mean squared error (RMSE), and mean absolute error (MAE). The detailed formulas and calculation methods for these evaluation metrics are as follows:

Mean Squared Error:

$$\mathrm{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{1}$$

MSE measures the average of the squared differences between predicted and true values, reflecting the sum of squared prediction errors.

Root Mean Squared Error:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{2}$$

RMSE, as the square root of MSE, provides an error metric in the same unit as the original data, making the results more interpretable.

Mean Absolute Error:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{3}$$

MAE calculates the average of the absolute differences between predicted and true values, offering the average magnitude of prediction errors and being less sensitive to outliers compared to MSE.

4.3. Analysis of model solving

The LSTM model consists of two primary layers. The first is an LSTM layer, and the second is a dense layer, which primarily maps the high-dimensional features output by the LSTM layer to the final predicted values, i.e., specific stock price numbers. The model is compiled using the Adam optimizer and MSE as the loss function. The Adam optimizer, an adaptive learning rate-based optimization algorithm, automatically adjusts the learning rate during training, making it suitable for handling various complex optimization problems. As a loss function, MSE effectively measures the difference between predicted and true values, driving the model to continuously optimize weight parameters by minimizing this error and improving prediction accuracy.

During the model training stage, preprocessed training data are fed into the constructed LSTM model. After balancing training time and overfitting risk, the number of training epochs is set to 20, and the batch size is configured as 32—a size that strikes a good balance between ensuring model update frequency and fully utilizing computational resources. During training, the model inputs data into the network in batches: it calculates predicted values through forward propagation, computes the error between predictions and true values using the loss function, and updates the network's weight parameters via the backpropagation algorithm. This process iterates until all training epochs are completed.

As shown in Figure 9 of this experiment, the model's predicted values exhibit a high degree of fit with true values, though some differences remain. The evaluation indicator results are as follows: MSE is 0.408, RMSE is 0.638, and MAE is 0.380. These values indicate that the model has a certain level of accuracy and stability in stock price prediction. An MSE of 0.408 suggests that the average squared error between predicted and true values is moderate: the model can fit the data reasonably well but still has room for improvement. An RMSE of 0.638 intuitively reflects that the average difference between predicted and true values is approximately 0.638 units, which is a reasonable range for stock price prediction and indicates the model can capture price fluctuations to some extent. A MAE of 0.380 demonstrates that the model maintains good stability in predicting daily price fluctuations, with relatively small errors, providing the investors with reliable references. Overall, the model performs acceptably in terms of prediction accuracy and stability, but it requires further optimization and improvement to handle extreme scenarios such as significant price fluctuations, thereby enhancing its adaptability and robustness.

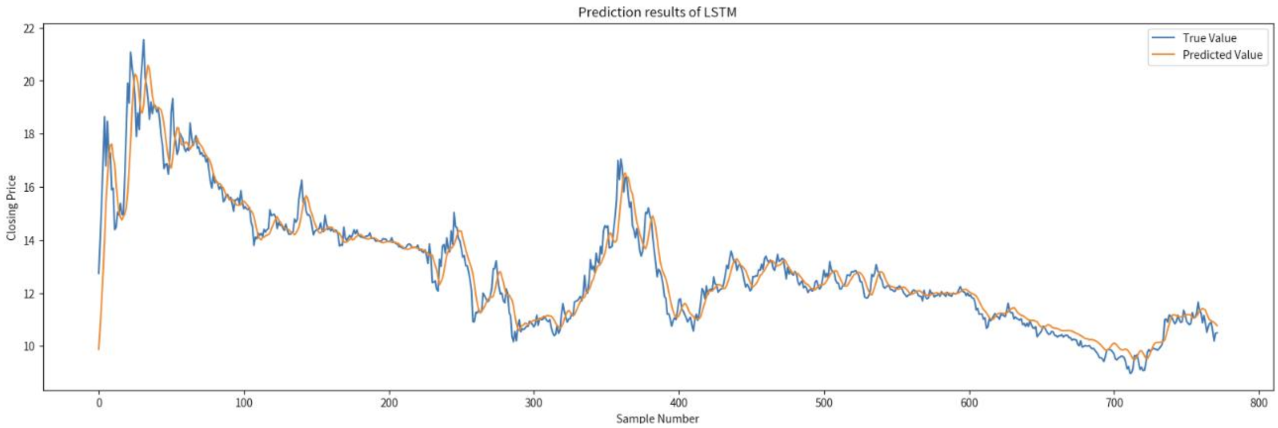

Figure 9. Prediction results of LSTM

4.4. Model Optimization

To enhance the LSTM model's performance in stock price prediction, this study conducted systematic hyperparameter optimization. Hyperparameters significantly influence a model's fitting capability and generalization performance; thus, optimizing them to find the optimal model configuration is critical. An exhaustive search method was employed for hyperparameter optimization, which identifies the best-performing combination on the validation set by iterating through all possible hyperparameter combinations. Although computationally expensive, this approach is effective when the search space is relatively small.

During the hyperparameter search, a reasonable range of search spaces was defined, including the number of LSTM layer neurons, dropout rate, optimizer type, batch size, and training epochs. For each hyperparameter combination, a model was created, trained, and evaluated on the training set using MSE as the scoring criterion. By traversing all combinations, the parameter settings that maximized model performance were ultimately identified. After completing the hyperparameter search, an optimized LSTM model was created using the found optimal parameter combination and retrained to ensure the model achieved its best performance.

As shown in Figure 10, the optimized model's predicted values (green line) exhibit a significantly closer fit to the true values (blue line) compared to the pre-optimization predictions (orange line). The optimized model performs better in capturing both the overall stock price trend and short-term fluctuations, particularly at key peak and trough positions where the gap between predicted and true values is notably reduced. This indicates that hyperparameter optimization effectively improves the model's prediction accuracy and stability. The evaluation indicators of the optimized model show a MSE of 0.172, RMSE of 0.415, and MAE of 0.294—all improved compared to pre-optimization values. Future research could consider adopting more efficient optimization algorithms, such as Bayesian optimization, to further enhance optimization efficiency and model performance. Additionally, incorporating more market data and technical indicators to enrich the model's input features could further strengthen its predictive capabilities.

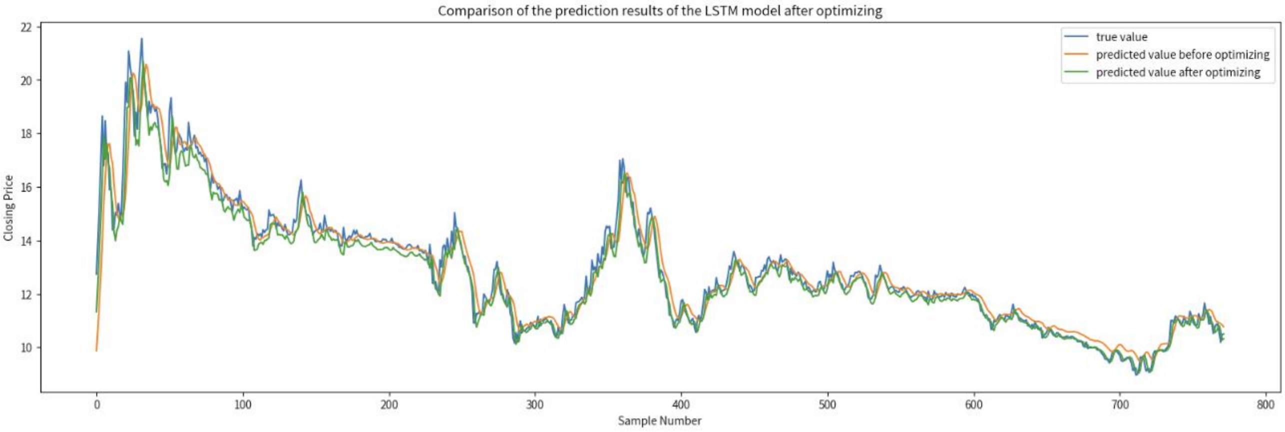

Figure 10. Comparison of the prediction results of the LSTM model

4.5. Conclusion of the experiment

Following data preprocessing and indicator construction, the model demonstrated strong learning capabilities during training, effectively fitting the training data. Hyperparameter optimization significantly enhanced model performance: by employing an exhaustive search method to optimize the LSTM model's hyperparameters, the optimal parameter combination was identified, minimizing the MSE on the validation set. The optimized model showed improvements in both prediction accuracy and stability, with evaluation metrics such as MSE, RMSE, and MAE all improving—indicating that hyperparameter optimization effectively strengthened the model's generalization ability and predictive performance.

The optimized LSTM model can predict stock price trends relatively accurately, providing valuable references for investors. The predicted values exhibit a high degree of consistency with the true values in terms of overall trends and short-term fluctuations, particularly at critical peak and trough positions where the fit between predictions and actual values is notable. This helps investors grasp market timing and formulate rational investment strategies.

5. Conclusion

This study successfully constructed and optimized an LSTM model, enabling it to demonstrate excellent performance in stock price prediction tasks. In the future, further exploration of more efficient hyperparameter optimization algorithms can be conducted to reduce computational costs and enhance optimization effectiveness. Additionally, incorporating more market data and technical indicators to further enrich the model's input features is expected to further improve its predictive capabilities, providing stronger support for analysis and decision-making in financial markets.

References

[1] Mo Jiawei. Research on Stock Price Prediction and Stock Selection Strategy Based on LSTM Model [D]. Guangzhou University, 2023.

[2] Yang Qing, Wang Chenwei. A Study on Global Stock Index Prediction Using Deep Learning LSTM Neural Network [J]. Statistical Research, 2019, 36 (03): 65–77.

[3] Cui Jianmin. Research on Stock Prediction and Application of Siamese Network Model Based on LSTM [D]. University of International Business and Economics, 2022.

[4] Dou Wei. Stock Price Trend Prediction Based on LSTM [D]. Tianjin University, 2020.

[5] Ren Jun, Wang Jianhua, Wang Chuanmei, et al. Stock Index Prediction Based on Regularized LSTM Model [J]. Computer Applications and Software, 2018, 35 (04): 44–48+108.

[6] Song Gang, Zhang Yunfeng, Bao Fangxun, et al. Stock Prediction Model Based on Particle Swarm Optimization LSTM [J]. Journal of Beijing University of Aeronautics and Astronautics, 2019, 45 (12): 2533–2542.

[7] Peng Yan, Liu Yuhong, Zhang Rongfen. Modeling and Analysis of Stock Price Prediction Based on LSTM [J]. Computer Engineering and Applications, 2019, 55 (11): 209–212.

[8] Zhao Hongrui, Xue Lei. Research on Stock Prediction Based on LSTM-CNN-CBAM Model [J]. Computer Engineering and Applications, 2021, 57 (03): 203–207.

[9] Hu Yamei. Stock Price Prediction Based on GARCH Family Models and LSTM Model [C]. School of Statistics and Mathematics, Guangdong University of Finance and Economics, 2023: 4.

[10] Liao Chang. Design and Implementation of a Stock Prediction Model Based on News Sentiment Quantification and LSTM Network [D]. Huazhong University of Science and Technology, 2020.

[11] Yu Xiaojian, Liu Guopeng, Liu Jianlin, et al. Stock Index Prediction Based on LSTM Network and Text Sentiment Analysis [J]. Chinese Journal of Management Science, 2024, 32 (08): 25–35.