Application of Deep Learning Methods in Predicting Stock Market Direction

Zhengyu Ding

{ dingzhengyu2005@outlook.com}

School of Economics, Huazhong University of Science and Technology, Wuhan, 430000, China

Abstract. This paper selects daily data from the S&P 500 from 2013 to 2024 and constructs a binary prediction box for stock market direction using average market return as the core indicator. Logistic regression is used as the baseline method to characterize linear relationships; Long Short-Term Memory (LSTM) is used to model the dynamic dependency structure in time series. Empirical results show that LSTM has a higher prediction accuracy than logistic regression, but the AUC values of both models are close to stochastic levels, indicating that it is difficult to judge market direction based solely on historical price information. It also shows that although deep learning models have advantages in representing nonlinear features, the improvement in predictive ability in noisy financial environments remains limited. This study provides empirical evidence for the selection of methods for predicting overall market trends.

Keywords: Stock market forecasting, Average market return, Logistic regression, LSTM, Time series modeling

1 Introduction

The stock market is undoubtedly a crucial part of the modern financial system, and the vast majority of investment decisions are based on its volatility. Due to its high volatility and non-stationary dynamics, the stock market presents substantial challenges for predictive modeling. Traditional stock prediction methods mostly use statistical models such as linear regression. These methods rely on strong assumptions of linearity or stationarity, which do not reflect the realities of the stock market, resulting in poor predictive performance. As deep learning continues to advance, Long Short-Term Memory (LSTM) networks have gained increasing attention in financial forecasting, Consequently, investigating the use of LSTM models in stock market prediction has gained growing attention in recent years.

Campisi et al. compared the performance of various machine learning models in predicting the direction of the US stock market [1]. The study found that some models outperformed the benchmark method in terms of accuracy, but the AUC value was close to the random level, reflecting that the market direction judgment itself has strong uncertainty. LSTM, as an important tool for time series modeling, is widely used for index trend prediction. Bhandari et

ICIAAI 2026, May 22-24, Singapore, Singapore

DOI 10.4108/eai.22-5-2026.2365120

al. used LSTM to model the stock index trend [2], demonstrating the application potential of deep learning in characterizing nonlinear financial sequences. Zhang and Pinsky compared the price change predictions of S&P500 and Nasdaq-100 [3], and believed that the prediction results at the index level were more stable than those at the individual stock level. Kundu and Pinsky pointed out in their daily direction prediction research that deep learning models usually outperform traditional models in terms of accuracy [4], but the improvement is limited in high-noise environments. Sang and Li introduced an attention mechanism into the LSTM structure to enhance the ability to identify key time features and improve model performance [5].

Traditional methods still have practical significance in financial classification tasks. Baguda et al. used price-perceived logistic regression to classify portfolio returns [6], and the results showed that logistic regression still has robust benchmark value. Yang et al. used group-penalized logistic regression to predict stock price trends [7], and achieved relatively stable prediction results after introducing feature selection and regularization. Thakkar and Chaudhari pointed out that financial market prediction faces problems such as non-stationarity [8], data leakage and overfitting, which has a direct impact on the model's ability to generalize to unseen data. Su et al. constructed a multi-factor LSTM model and integrated multi-dimensional features into the time series framework to improve prediction performance [9]. Sherstinsky systematically explained the gating mechanism and gradient propagation process of RNN and LSTM [10], providing theoretical support for the application of deep models in time series tasks.

This paper applies both a traditional logistic regression model and a deep learning approach, namely LSTM, to predict the overall direction of the stock market, and compares the predictive abilities of the two models to demonstrate the effectiveness of deep learning models that can capture long-term time-dependent features in stock market direction prediction. Meanwhile, this paper uses an 11-year long-term, large-sample dataset to effectively avoid the influence of random factors on model evaluation, providing valuable experience for future research on stock market prediction.

2 Research methods

2.1 Introduction to the dataset

This article will use a 2013-2024 dataset of stock information of S&P 500. It consists of the opening price, closing price, highest price and lowest price of all the stocks of S&P 500 and it has a total of 1, 048, 576 records. The source of data is Kaggle (S&P 500 Stocks (daily updated)). The credibility score against Kaggle on the dataset is 10, which is valid as far as data science research goes. In the present paper, the dataset is partitioned into training and test set in a balanced 4:1 proportion on time series data.

2.2 Problem modeling

As everyone knows, the stock market can be characterized as very random, and the prices of individual stocks are very vulnerable to random shocks, and the individual stock prices highly fluctuate to cause dramatic changes which could not be explained by the logic of prediction. Hence, it is very hard to predict individual stocks. Contrastingly, an increase in the research object of individual stocks into the whole stock market can be effective in minimizing the effect of random shocks on individual stocks hence a clearer capture and prediction of the overall stock market trend will be possible. Hence, it is upon this paper to make the decision to model the problem in the overall market level. This paper in particular employs the average market return (cross-sectional average of all stocks in a given trading day) as a measure to describe the overall performance and indicate the general movement of the market on that particular trading day.

As a binary classification problem, in line with the average market return, this paper converts this into the market prediction. When the average market return in the coming trading day is positive, then then market is considered to be rising, represented as category 1, otherwise the market is considered to be neither rising nor falling which is represented as category 0. This eases the adoption of a classification model to forecast the general market trend.

2.3 Baseline model: logistic regression

The Logistic Regression is an old type of binary classification that has been heavily utilized in other areas of research like financial risk assessment and market direction prediction. The baseline model applied in the research structure of this paper is the logistic regression model. Where t the market feature vector at time t is:

$$\mathbf{x}_t = (x_{t,1}, x_{t,2}, \dots, x_{t,p})^\top \tag{1}$$

Wherein $p$represents feature dimension. First, the features are linearly weighted to construct a linear prediction function:

$$z_t = \beta_0 + \sum_{j=1}^{p} \beta_j x_{t,j} \tag{2}$$

Where is β_0 the intercept term, β_j and is the regression coefficient of the corresponding feature.

Subsequently, the linear prediction result is mapped to the probability of a market rise using the Sigmoid function:

$$P(y_{t+1} = 1 \mid \mathbf{x}_t) = \sigma(z_t) = 1/(1 + e^{-z_t}) \tag{3}$$

Among them, $y_{t+1} \in \{0,1\}$ the category label indicates the overall market trend on the next trading day, which $y_{t+1} = 1$ $y_{t+1} = 0$ indicates that the market is rising, falling, or not rising.

The model parameters are solved using maximum likelihood estimation, and the corresponding loss function is the negative of the log-likelihood function, which can be expressed as:

$$L(\beta) = -\sum_{t=1}^{T} \left[y_{t+1} \log P(y_{t+1} = 1 \mid x_t) + (1 - y_{t+1}) \log(1 - P(y_{t+1} = 1 \mid x_t)) \right] \tag{4}$$

Where is T the sample time length.

In actual prediction, when the probability of an upward trend output by the model is greater than a preset threshold (set to 0.5 in this paper), the market is judged to be in an upward trend; otherwise, it is judged to be in a downward (or balanced) trend. In this way, the logistic regression model can map continuous market characteristics into clear directional judgment results.

2.4 LSTM model construction

Recurrent neural networks are used to predict time series and one of them is the Long Short-Term Memory (LSTM) network. Existing densely connected recurrent neural networks tend to suffer a vanishing gradient, or gradient explosion, when dealing with longer time series, and LSTM alleviates these effects by implementing a gating mechanism which enables the model to learn both short term variations and long term dynamics.

A single-layer LSTM network having 64 hidden units and the time window length of 20 are used in this paper. In order to reduce overfitting, there is the application of Dropout = 0.2 between the LSTM and fully connected layers. The task of the classification head includes 32-dimensional ReLU fully connected layer, and the Sigmoid output layer. The paper experiments with the Adam optimizer and binary cross-entropy loss and utilizes Early Stopping to automatically terminate training based on the accuracy on validation set, balancing the capacity of the model to fit and the generalization of the model.

The main concept of LSTM is the addition of three gating networks, namely, input gate (see Equation (6)), forget gate (see Equation (5)) and output gate (see Equation (7)). The forget gate controls which elements of the historical information are to be remembered, the input gate controls the amount of the new information to be written at the present point in time and the output gate controls the contribution that is to be made by current hidden state to the external output. In the above gating process, LSTM can dynamically choose significant information and store it hence adjusting to long-term trend learning.

$$f_t = \sigma(W_f \cdot [h_t - 1, x_t] + b_f) \tag{5}$$

$$i_t = \sigma(W_i \cdot [h_{\{t-1\}}, x_t] + b_i) \tag{6}$$

$$o_t = \sigma(W_o \cdot [h_{\{t-1\}}, x_t] + b_o) \tag{7}$$

In the model design presented in this paper, LSTM is used to characterize the dynamic features of the overall market indicators over time. The input layer receives time-series data containing historical market characteristics, then feeds the input sequence into the hidden layer, updates the historical information through a gating mechanism, and finally the output layer generates a prediction result for the next trading day based on the output of the hidden layer.

With respect to the design of the output layer, this paper uses the Sigmoid function to map the model output to the probability of a market rise. When the predicted probability exceeds a preset threshold, the model determines that the market is in an upward trend; otherwise, it determines that the market is in a downward trend or not rising.

$$P(y\_(t+1) = 1 \mid \mathbf{x}\_t) = \sigma(z\_t) = 1/(1 + e\text{^}(-z\_t\ )\ ) \quad (8)$$

The LSTM model was chosen for this study primarily based on the following considerations. First, financial market time series typically exhibit significant time dependencies, and LSTM has a natural advantage in modeling such time dependencies. Second, LSTM has been validated in numerous financial time series forecasting studies, and as a mature deep learning model, it is suitable as the main nonlinear forecasting method used in this paper to characterize the overall market trend.

2.5 Evaluation indicators

To comprehensively evaluate the model's performance in market forecasting tasks, this paper uses a variety of commonly used classification evaluation metrics to compare the model's performance.

Accuracy measures the proportion of model predictions that match the true labels. Essentially, it is the sum of samples that correctly predicted an increase and those that correctly predicted a decrease (or no increase) divided by the aggregate number of samples. The metric can be used to measure the accuracy of the model's predictions.

For the further characterizing that the model's ability to identify upward trends, this paper introduces precision and recall. Precision represents the proportion of samples predicted as upward trends that actually are; recall represents the proportion of samples that are actually upward trends that the model successfully identifies as upward trends. The F1 score is the harmonic mean of precision and recall, used to comprehensively measure the model's predictive performance in the upward trend direction.

Based on the above indicators, this paper evaluates different models from two aspects: overall accuracy and the ability to identify the types of upward trends, thereby comparing the performance of logistic regression and LSTM in market trend prediction.

3 Results

The logistic regression model did the test set with an accuracy of 52.53 which is slightly more than random guessing as seen in Table 1. At the same time, the model obtained a recall of 95.43% when the market trend was upwards, which means that it successfully classified the presence of the real upward market situation. The model was however limited to forecasting down market conditions or stagnant market conditions which was weak. This indicates that simple linear models are more effective in forecasting the general trend of the market yet fails to distinguish the various states of the market.

On overall market time series data, confidence of the LSTM model in predicting data was 54.99, which is much better compared to the logistic regression model. This shows that with time series modeling, LSTM can be more useful to hold the dynamic nature of how market state evolves over time. The recall of LSTM was however only 92.87, which means that it suffers the same issue as the logistic regression model, namely it is not powerful enough to distinguish different state of the markets.

Both Figure 1 show that the ROC curves of the two models are near to a diagonal distribution and the AUC value is relatively similar to the baseline level of 0.5. This indicates that the predictive accuracy of the two models is not particularly high and it also reveals that when predicting the right direction of the stock market with a high amount of noise then predicting with single-indicator modeling is hard.

LSTM models compared with logistic regression models have a greater ability to express nonlinearities in the explanation of the overall market trends because they take historical time-series data into account. This outcome confirms the arguments posing the deep learning techniques into the market level prediction tasks.

Table 1. Model Performance Comparison.

Index	Accuracy	Precision	Recall	F1 score
Logistic Regression	0.52527	0.5271 1	0.95431	0.67910
LSTM	0.54992	0.54624	0.92878	0.68791

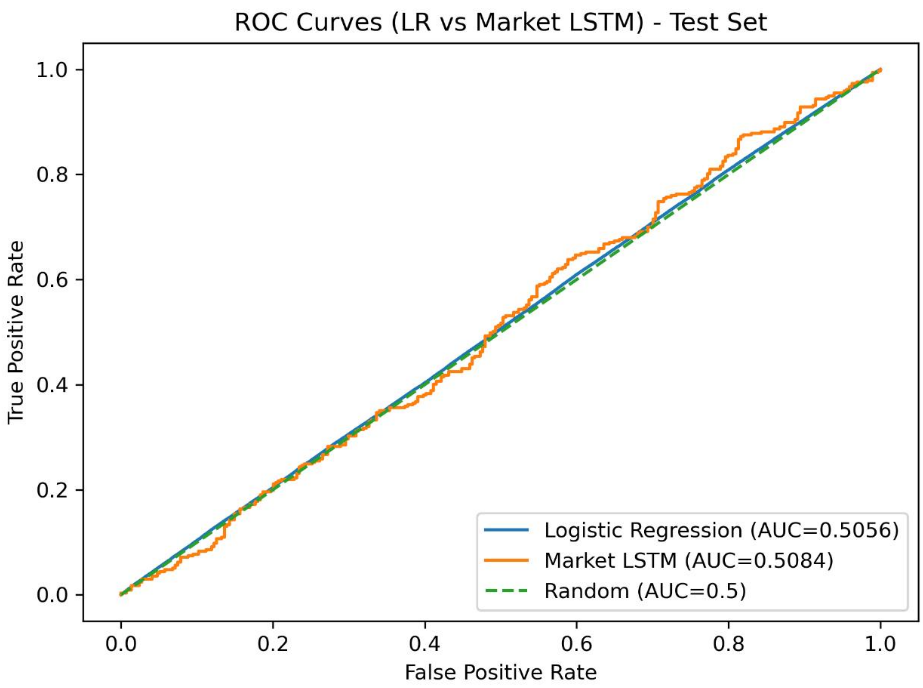

Fig. 1. ROC curves of the two models. (Picture credit: Original)

4 Conclusion

To predict the direction of the market, the given paper will apply a binary classification prediction model suggested by average market returns and utilize both logistic regression and LSTM to investigate the market direction and predict it. Comparison and analysis of the findings of various types of models have helped the paper to come up with the following conclusions:

To start with, concerning the prediction result, a logistic regression model is the simplest model, with high interpretability and stability, although its predictive potential is quite low when applied to time series data, at the same time, the LSTM model has a time model mechanism and is more successful in prediction accuracy. Meanwhile, the two models are more effective compared to random prediction in the predictive tasks, meaning it is possible to predict the general direction of the market under consideration on the basis of the historical data.

Second, based on the performance disparities among the various models, logistic regression model faces some limitations on its time series prediction facilities since it fails to identify accurately the market situations and hence makes error in predictions. Although predicting the markets is also problematic with the LSTM model, there are more obvious issues with it predicting non-linear relations and more effectively responding to the time series which leads to a better predictive capabilities. Nonetheless, due to the nature of both stock market and their associated noises and uncertainties, they cannot predict to a great degree, and it is therefore relatively limiting in predictive effect. The task of predicting a financial market is not easy. In spite of the fact that application of average market returns is able to reduce non-systematic noises, at least in individual persons of stock, numerous aspects, like policy alterations, unforeseen events, and abrupt favorable news, nonetheless play a major role in a market performance, and they are not simple to reflect and reduce with the support of pure market statistics. Future studies might consider adding datasets like sentiment analysis and trying to turn policy benefits into something that is modelable so as to increase predictive abilities. Moreover, the use of more integrative points of measurement, along with time-based prediction models might also make a big difference in predictive power.

References

[1] Campisi, G., Muzzioli, S., De Baets, B.: A comparison of machine learning methods for predicting the direction of the US stock market on the basis of volatility indices. International Journal of Forecasting. (2023)

[2] Bhandari, H. N., et al.: Predicting stock market index using LSTM. Machine Learning with Applications. (2022)

[3] Zhang, X., Pinsky, E.: S&P-500 vs. Nasdaq-100 price movement prediction with LSTM for different daily periods. Machine Learning with Applications. (2024)

[4] Kundu, T., Pinsky, E.: Predicting daily stock price directions with deep learning models. Machine Learning with Applications. (2025)

[5] Sang, S., Li, L.: A Novel Variant of LSTM Stock Prediction Method Incorporating Attention Mechanism. Mathematics (MDPI). pp. 945 (2024)

[6] Baguda, Y. S., AlJahdali, H. M., Taha, A. A.: Dynamic Portfolio Return Classification Using Price-Aware Logistic Regression. Mathematics (MDPI). pp. 1885 (2025)

[7] Yang, Y., Hu, X., Jiang, H.: Group penalized logistic regressions predict up and down trends for stock prices. The North American Journal of Economics and Finance. pp. 101564 (2022)
[8] Thakkar, A., Chaudhari, K.: A comprehensive survey on deep neural networks for stock market: the need, challenges, and future directions. Expert Systems with Applications. pp. 114800 (2021)
[9] Su, Z., Xie, H., Han, L.: Multi-factor RFG-LSTM algorithm for stock sequence predicting. Computational Economics. pp. 1041–1058 (2021)
[10] Sherstinsky, A.: Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. Physica D: Nonlinear Phenomena. pp. 132306 (2020)