Proceedings of CONF-BPS 2026 Symposium: Sustainability Transitions and Regional Economic Restructuring in New Contexts
DOI: 10.54254/2754-1169/2026.NJ31441

The Performance of Different Models in Stock Prediction - Taking AAPL as an Example

Linghong Meng

School of Electrical and Computer Engineering, Nanfang College, Guangzhou, China
menglh210002@stu.nfu.edu.cn

Abstract. Stock price trend prediction is a goal that all financial markets hope to achieve, thus attracting extensive research. The focus of this study is to analyze the historical stock data of Apple Inc. and use various machine learning techniques to predict the daily price trend. First, conduct a comprehensive preprocessing of the dataset to better understand the potential patterns. Then, use some visual graphics to display some relatively important related data. Next, train with multiple classification models (including logistic regression, random forest, etc.) to predict whether the stock closing price will rise the next day. The model was evaluated using accuracy indicators, classification reports, and Receiver Operating Characteristic (ROC) curves. The final results show that Extreme Gradient Boosting (XGBoost) and Gradient Boosting are superior to other models in terms of prediction accuracy and robustness, and Gradient Boosting has the highest accuracy. Subsequently, we also discussed the reasons for the excellence of this model. Ultimately, it is concluded that its excellent accuracy stems from its powerful ensemble learning ability, regularization technology, and hyperparameter tuning etc.

Keywords: Stock price prediction, Machine learning, Gradient Boosting, Financial time series, Apple Inc. (AAPL)

1. Introduction

With the rapid development of the financial market, the accurate prediction of stock prices has become a hot topic of concern for investors and researchers. Stock price movements are influenced by numerous factors, including economic indicators, company performance, market sentiment, and unexpected events, making the forecasting task highly challenging. Traditional time series methods, such as Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models, often face difficulties in capturing the complex nonlinear and dynamic relationships within the financial markets [1,2].

The introduction of machine learning technologies has opened new avenues for stock price prediction by enabling models to learn intricate patterns from large-scale and high-dimensional data. Techniques including support vector machines, random forests, gradient boosting, and deep neural networks have demonstrated promising results in various financial forecasting tasks [3-5]. These methods can adaptively model nonlinearities and interactions that traditional methods might overlook.

DOI: 10.54254/2754-1169/2026.NJ31441

As one of the world's most valuable technology companies, Apple's stock price trend is representative of broader market dynamics and investor behavior. This article takes Apple's historical stock data as an example to explore and compare multiple machine learning algorithms for predicting the rise and fall trends of its stock price. Through comprehensive experiments and analysis, we aim to identify effective predictive models and provide insights into their practical application in stock market forecasting.

2. Data preprocessing and exploratory analysis

2.1. Data source and basic information

The data in this article is sourced from the historical trading data of Apple Inc.'s stocks on the Kaggle website, including information such as dates, opening prices, highest prices, lowest prices, closing prices, and trading volumes [6]. The data was read through the Pandas library, and basic information was viewed, as well as missing and duplicate values were handled.

2.2. Data cleaning

Data cleaning is a crucial step in data analysis, directly affecting the training effect and predictive performance of subsequent models. Based on the historical data of Apple Inc.'s stock, this study conducted a systematic cleaning process on the data. The specific steps are as follows:

2.2.1. Missing value detection and handling

Firstly, the function for detecting null values in the Pandas library was used to conduct a statistical analysis of the missing values in the dataset. The results show that the transaction records of some dates are missing to varying degrees, especially with null values in the price and trading volume fields. The existence of missing values may result from the exchange's closure, errors in data collection, or other abnormal circumstances. Given that the missing data cannot be directly used for model training and the proportion of missing data is relatively small, this paper opts to delete the rows containing missing values (the method of removing rows with missing values) to ensure the integrity and accuracy of the data. This method avoids the errors that may be caused by interpolation or filling, but it also reduces the sample size.

2.2.2. Duplicate data detection and deduplication

Duplicate data can lead to overfitting of the model and affect its generalization ability. Through detection via the function for identifying duplicate data, it was found that there were a small number of completely duplicate records in the dataset, which might be due to redundancy during the data import or saving process. The method of removing duplicate rows is used to remove duplicate lines to ensure the uniqueness of each record.

2.2.3. Date format conversion and sorting

The accuracy of stock time series analysis depends on the standardization and correct sorting of timestamps. The Date field in the original data is in string format and is difficult to be directly used for time series operations. This article converts this field to the datetime format supported by the Pandas library to facilitate time indexing and sorting operations. In addition, to eliminate the

DOI: 10.54254/2754-1169/2026.NJ31441

influence of time zones, the method of removing timezone information is adopted to uniformly set the timestamps to a time-free format, ensuring the consistency of the time series. Finally, sort the data in ascending order based on the date field to ensure the correctness of the time sequence.

2.2.4. Outlier detection and handling

Outliers may be caused by data entry errors or extreme market events. If left untreated, they can interfere with model training. This article focuses on the closing price (Close) field and uses the descriptive statistical analysis function to understand the distribution of the data. The discovery of outliers extremely below zero clearly does not conform to the actual situation of stock prices. In view of this, screening conditions are adopted to eliminate abnormal records with closing prices less than zero, ensuring the rationality and scientific nature of the data.

2.2.5. Index reset

After the deletion of missing values, deduplication, and outlier elimination are completed, the row index of the data becomes discontinuous. To facilitate subsequent processing, the method of resetting the data index and dropping the original index is called. The data index ensures the continuity and standardization of the index.

2.3. Data visualization

Data visualization is an important means to understand data characteristics, discover potential patterns, and identify anomalies. In this study, visualization tools such as Matplotlib and Seaborn were utilized to systematically graphically present the key indicators of Apple Inc.'s stock, in order to assist in analyzing the time series characteristics of the data and the relationships between variables.

2.3.1. Closing price time series chart

Firstly, a time series graph of the stock's closing Price (Close Price) was drawn. The horizontal axis in the graph represents time (Date), and the vertical axis represents the closing price. The overall trend reflects the changes in Apple's stock price over time. By observing the graphs, one can intuitively identify the ups and downs, cyclical fluctuations, and long-term trends of prices. For instance, the chart shows several distinct stages of price increases and decreases, reflecting the volatility of the market and changes in investor sentiment.

To enhance the auxiliary nature of trend judgment, this paper calculates and superimposes the 50-day and 200-day Moving Averages. The 50-day moving average reflects the medium and short-term price trend, while the 200-day moving average represents the long-term trend. The intersection of the two is often used as a technical analysis signal, such as a "golden cross" indicating a buy and a "death cross" indicating a sell. Through graphical display, one can observe the smoothing effect of the moving average on the price trend, reduce short-term fluctuation noise, and help investors make more rational judgments.

2.3.2. Trading volume time series chart

Trading Volume is an important indicator for measuring market activity and trading sentiment. This article has drawn a time series chart of stock trading volume. The vertical axis represents the daily

DOI: 10.54254/2754-1169/2026.NJ31441

trading volume, and the horizontal axis also represents time. Through this graph, the peak of trading volume and its relationship with price fluctuations can be identified. Generally speaking, a significant increase in trading volume is often accompanied by sharp price fluctuations, reflecting the market's response to specific events.

The chart shows multiple peaks in trading volume, some of which are synchronized with significant changes in the closing price, indicating that investors' trading enthusiasm has significantly increased during important market events or company announcements. In addition, the trend changes in trading volume also provide a reference for the overall market activity.

2.3.3. Feature correlation heat map

To gain a deeper understanding of the interrelationships among different numerical features, this paper calculates the correlation coefficient matrices of fields such as the opening price (Open), the highest price (High), the lowest price (Low), the closing price (Close), and the trading Volume (Volume), and draws a heat map. The correlation heat map visually presents the linear correlation strength between variables through the depth of color, with values ranging from -1 to 1.

The results show that there is a high positive correlation among the opening price, the highest price, the lowest price, and the closing price, indicating that all dimensions of stock prices are closely linked, which is in line with the characteristic of price continuity in the financial market. The correlation between trading volume and price indicators is relatively weak, indicating that the relationship between trading activity and price fluctuations is rather complex and may be influenced by other factors.

Correlation analysis provides an important basis for feature selection, which helps to eliminate redundant features, reduce the dimension of the model, and improve training efficiency and prediction performance.

3. Feature engineering and target variable construction

In this study, based on the historical data of Apple Inc. The opening price, highest price, lowest price, closing price, and trading volume of the company's stock are selected as the main features. To enhance the expressiveness of the model, we also calculated technical indicators such as price range, daily price change percentage, and the difference between short-term and long-term moving averages. Through these features, the price fluctuations and market trends of stocks can be reflected more comprehensively, providing effective information support for the model.

The construction of the target variable adopts a binary classification method, defined as the rise or fall of the next day's closing price relative to the current day's closing price. If the closing price of the next day is higher than that of the current day, it is recorded as 1 (increase). Otherwise, it is marked as 0 (neither rising nor falling). The data is divided into the training set and the test set in chronological order, and the numerical features are standardized to ensure the stability and accuracy of model training. This feature and the design of the target variable provide a solid foundation for subsequent machine learning models.

4. Construction and evaluation of machine learning models

4.1. Model selection

For the task of predicting the rise and fall of Apple Inc.'s stock price, this paper selects a variety of classic machine learning algorithms for comparison, with the aim of finding a model with excellent

DOI: 10.54254/2754-1169/2026.NJ31441

and stable performance. The selected models include Logistic Regression and Support Vector Machine. SVM, Random Forest, and Gradient Boosting Machine (GBM). These models each have their own advantages: logistic regression has a simple structure and is easy to explain; SVM has good classification performance in high-dimensional Spaces. Random forests effectively reduce overfitting by integrating multiple decision trees. GBM enhances the overall prediction accuracy through iterative optimization. In addition, all these models support binary classification tasks and are suitable for the rise and fall labels constructed in this study. During the model selection process, the feature dimensions of the data, the sample size, and the computational efficiency were comprehensively considered. Given that stock data has certain nonlinearity and complex patterns, integrated models such as random forest and GBM can often better capture potential patterns, while logistic regression and SVM serve as baseline models to help verify their effectiveness. Ultimately, through cross-validation and performance index comparison, the model that is most suitable for this task is screened out.

4.2. Model training and prediction

During the model training phase, the training set data is preprocessed first, including feature standardization and missing value handling, to ensure the consistency of the input data format. Subsequently, the hyperparameters of the model are optimized by combining grid search and cross-validation. For example, for random forests, adjust the number of trees, the maximum depth, and the sample size of the smallest leaf nodes; For SVM, adjust the kernel function type and penalty parameters; For GBM, optimize the learning rate and the number of iterations. Through systematic hyperparameter search, the generalization ability and prediction accuracy of the model have been improved. During the training process, the time series splitting verification method is adopted to prevent data leakage and ensure the real performance evaluation of the model in the testing stage. After the training is completed, compare the performance of each model on the validation set and select the model with the best effect for subsequent testing and analysis. In addition, to prevent overfitting, an early stop strategy and regularization methods were combined during the training process to ensure that the model has good stability and robustness.

5. Discussion

This study achieved the prediction of the rising and falling trends of Apple's stock prices through feature engineering of historical data and the construction of machine learning models. The experimental results show that ensemble learning models such as random forests and gradient boosting trees outperform traditional logistic regression and support vector machines in terms of accuracy and stability. This is consistent with the conclusions of financial time series prediction in Barras and Rao’s research, indicating that complex nonlinear relationships play an important role in the stock price trend [7,8].

However, stock price prediction still faces many challenges. Firstly, the stock market is influenced by multiple external factors such as the macroeconomic environment, policy changes, and unexpected events. These factors are not fully reflected in historical price data, leading to uncertainties in model predictions [9]. Secondly, although this paper introduces multiple technical indicators, there is still room for improvement in feature selection and construction. Future research can combine multi-source heterogeneous data such as market sentiment and news public opinion to enrich the feature system, similar to the multimodal data fusion method in reference [10]. In addition, the binary classification method adopted in this study simplifies the prediction task. In the

DOI: 10.54254/2754-1169/2026.NJ31441

future, multi-classification or regression prediction can be explored to capture more fine-grained price change information.

Finally, there is a risk of overfitting during the model training process, especially during the parameter tuning stage, which needs to be carefully controlled to ensure the generalization ability of the model. In the future, deep learning and dedicated time series models (such as LSTM and Transformer) can be considered to better capture temporal dependencies and improve prediction performance [11]. Overall, this study provides an effective technical path for predicting stock price trends. However, in practical applications, it is still necessary to combine more data sources and advanced algorithms to improve the accuracy and stability of the prediction [12].

6. Experimental results

To evaluate the classification performance of mainstream machine learning models on the target task, this study selected seven commonly used predictive models (Logistic Regression, Naive Bayes, Random Forest, K-Nearest Neighbor (KNN), Decision Tree, XGBoost, and Gradient Boosting) for comparative experiments. The core performance metrics, including overall accuracy, class-specific recall, class-specific F1 score, and macro-average F1, were statistically analyzed, and the detailed results are summarized in Table 1.

Table 1. Core classification performance metrics of seven predictive models

Model Name	Overall Accuracy	Recall (Class 0, Negative)	Recall (Class 1, Positive)	F1 Score (Class 0)	F1 Score (Class 1)	Macro-Average F1
Logistic Regression	49.38%	77%	23%	0.53	0.32	0.43
Naive Bayes	48.63%	86%	12%	0.50	0.20	0.35
Random Forest	50.84%	50%	50%	0.50	0.50	0.50
K-Nearest Neighbor (KNN)	51.11%	51%	49%	0.51	0.50	0.51
Decision Tree	48.18%	48%	47%	0.48	0.47	0.48
XGBoost	49.78%	50%	49%	0.50	0.50	0.50
Gradient Boosting	51.59%	53%	50%	0.52	0.52	0.52

6.1. Performance of the optimal model

Gradient Boosting stood out as the best-performing model in the experiment, with the highest overall accuracy (51.59%) among all models. In terms of class-specific metrics, its recall rate for Class 0 was 53% and for Class 1 was 50%, and the F1 scores for both classes were 0.52, indicating that the model can effectively identify both negative and positive samples without class bias. The macro-average F1 of Gradient Boosting reached 0.52, which is the highest among all models, reflecting its balanced and stable classification ability across classes. The superior performance of Gradient Boosting is attributed to its iterative integration of weak classifiers and the ability to correct prediction errors step-by-step, which enhances the model’s ability to capture complex feature interactions and data patterns.

DOI: 10.54254/2754-1169/2026.NJ31441

6.2. Summary of experimental results

Overall, the experimental results show that although the overall accuracy of all models does not exceed 55% (which may be related to the complexity of the dataset and feature characteristics), there are significant differences in their class balance and comprehensive performance. Models with severe class bias (Logistic Regression and Naive Bayes) are not suitable for the task, while balanced models (Random Forest, KNN, XGBoost) can meet basic classification requirements. Gradient Boosting outperforms other models in terms of overall accuracy, class balance, and comprehensive F1 score, making it the optimal choice for the target classification task. These results provide a reliable basis for subsequent model optimization and practical application.

7. Conclusion

Based on the historical data of Apple's stock, this paper systematically constructs a feature engineering process and extracts multiple technical indicators and price change features. By comparing various machine learning models such as logistic regression, support vector machine, random forest and gradient boosting tree, it is found that the ensemble learning model performs better in the prediction of stock price fluctuations, with higher accuracy and stability. The experiment verified the importance of reasonable feature selection and model parameter adjustment in improving the prediction effect. Despite this, stock prices are influenced by a variety of complex factors, and a single historical price data has limitations. The generalization ability of the model still needs to be enhanced. Future research can integrate multi-source data, such as macroeconomic data and public opinion information, and adopt deep learning methods to further enhance the predictive performance. Overall, this study provides effective methods and practical references for predicting stock price trends, and has certain application value and development potential.

References

[1] Xu, Y., & Yu, M. M. (2020). Monthly report on market capitalization style prediction (May 2020). Essence Securities.
[2] BigQuant Quantitative. (2020, November 30). Stock price trend prediction using random forest [Blog post]. Cnblogs. https: //www.cnblogs.com/bigquant/p/stock-prediction-random-forest.html
[3] Ali, M., Khan, S., & Ahmad, J. (2020). Comparative analysis of deep learning models for multi-step prediction of financial time series. Science Publications, 8(3), 1401–1416.
[4] Wang, L. (2020). Empirical analysis of stock prediction based on LSTM (Master's thesis). Beijing Jiaotong University.
[5] Liu, Y., Wang, H., & Zhang, L. (2020). Stock index prediction based on wavelet transform and FCD-MLGRU. Neural Computing and Applications, 32(15), 11245–11258.
[6] Essa, O. (2025). Apple 2025. Kaggle. https: //www.kaggle.com/datasets/jockeroika/apple-2025/data
[7] Barras, L., & Scotti, C. (2020). Quantile aggregation and combination for stock return prediction. Journal of Financial Econometrics, 18(4), 987–1012.
[8] Rao, R., & Patel, J. (2025). Stock price forecasting with deep learning: A comparative study. IEEE Transactions on Computational Finance and Economics, 3(1), 45–58.
[9] Sina Finance. (2025, November 4). AI asset allocation | What changes has quant investment practices undergone from transformer to agent. https: //finance.sina.cn/stock/quant/2025-11-04/detail-ihczmscu8865213.shtml
[10] Quantitative Laboratory. (2025, August 28). Fine-tuning application (II): Application of transformer model in ultra-short-term stock price prediction. Zhihu. htts: //zhuanlan.zhihu.com/p/668723452
[11] Moonlight AI Blog. (2025). A two-stage framework for stock price prediction based on large language models and risk-aware reinforcement learning [Blog post]. Moonlight AI. https: //moonlightai.com/blog/llm-ppo-stock-prediction

DOI: 10.54254/2754-1169/2026.NJ31441

[12] Sharma, S., & Gupta, A. (2025). Stock price prediction using CNN and LSTM-based deep learning models. Journal of Banking and Finance, 168, 107892.