

A Study on Stock Price Forecasting Based on LSTM and Cross-Modal Adaptation

Ronghua Zhang

Macau University of Science and Technology, Macao, Macao SAR, China

a479150660@gmail.com

ABSTRACT

Stock price forecasting is an important topic in financial engineering and intelligent computing. Traditional price-only models often fail to reflect market sentiment, public opinion shocks, and changes in investor expectations. To address this problem, this study uses three A-share stocks as cases: Ping An Bank (000001), BYD (002594), and Kweichow Moutai (600519). It builds a daily-level alignment framework for stock price data and stock forum comments. On the text side, the comments are cleaned, scored for sentiment, and aggregated by trading day. An influence-weighted sentiment representation is also constructed. At the model level, Logistic Regression, Price-only LSTM, Price + Sentiment LSTM, and Gating-LSTM are compared. The results show that sentiment features are difficult to use when price data and comments are not aligned by time. After the dataset was expanded to three stocks, the linear model and the standard LSTM still showed limited improvement from direct sentiment concatenation. The All-Stocks Gating-LSTM achieved a test accuracy of 0.7647. Its F1 score was 0.3333, and its mean gate value was 0.5958. These results suggest that the gating module participates in the prediction process. They also show that the gating mechanism changes the classification tendency of the model. Overall, the study supports the value of temporal alignment and provides a basis for further testing of text sentiment features and cross-modal gating in stock prediction.

Keywords: Stock Price Forecasting; LSTM; Sentiment Analysis; Multimodal Fusion; Gating Mechanism

1. INTRODUCTION

Stock market forecasting is an important research topic in financial engineering and intelligent computing. Its results can support investment decisions, risk management, and quantitative trading. Stock prices are noisy, non-linear, and non-stationary. They are also affected by historical market trends, macroeconomic conditions, unexpected events, and investor sentiment. Because of this, traditional linear models often perform poorly in complex market conditions. Existing financial NLP studies show that news, forum comments, and social media texts can be transformed into structured variables such as sentiment, attention, and topic factors [1-3]. Recent multimodal stock forecasting studies further suggest that combining price, text, and other heterogeneous information can provide additional predictive signals [4]. LSTM models are widely used in stock forecasting because they can model temporal dependencies in sequential data [5,6]. Hybrid LSTM systems that include sentiment information have also been shown to improve movement prediction in some settings [7]. However, several problems remain. Financial text is often noisy. Many methods still rely on static fusion. They also do not clearly model the changing importance of text across trading days. To address these issues, this study focuses on the A-share market. It constructs a daily-level alignment framework for price data and comment text. It also designs daily sentiment factors and influence-weighted sequences. Finally, it proposes a Gating-LSTM model. The model uses a daily-level gate to adjust the contribution of text features, following the general idea of gated multimodal fusion [8] and recent gated cross-modal stock prediction research [9].

2. LITERATURE REVIEW

2.1 Research on Stock Price Prediction Using Financial NLP and Sentiment Analysis

Internet platforms have produced a large amount of unstructured text in the stock market. Common examples include news articles, announcements, stock forum comments, and social media posts. The main task of financial NLP is to convert these texts into structured variables. These variables may include sentiment, attention, or topic information [1]. They are usually obtained through text pre-processing, sentiment analysis, and topic extraction. They can then be used for financial analysis and forecasting. Stock selection studies based on NLP also show that stock comments and news factors can support

investment-related prediction tasks^[2]. In addition, StockNet demonstrates that tweets and historical prices can be organized with trading-day alignment for stock movement prediction^[3]. Existing multi-source studies further show that investor sentiment can be combined with other market information to improve prediction tasks^[4]. However, financial text is noisy. This problem is more obvious on social media platforms. The text may contain colloquial expressions, repeated posts, and low-quality information. Simple equal-weighted aggregation can weaken useful signals. Therefore, robust daily sentiment factors are needed. It is also important to combine sentiment with interaction information through weighted aggregation.

2.2 LSTM-Based Stock Time Series Forecasting Research

LSTMs can reduce the vanishing gradient problem. They can also model long-term dependencies in sequential data. For this reason, they are widely used in stock time series forecasting. Fischer and Krauss show that LSTM networks can outperform memory-free classification methods in large-scale financial market prediction^[5]. Improved LSTM frameworks have also been proposed for stock prediction systems, including bidirectional structures, attention mechanisms, and residual connections^[6]. Recent hybrid forecasting systems further combine LSTM with sentiment analysis to use both historical prices and textual signals^[7]. However, much LSTM research still focuses mainly on structured price features. Even when text variables are added, they are often directly concatenated with price features. This approach does not explicitly model the dynamic contribution of each modality. In stock forecasting, the importance of text changes over time. As a result, standard LSTMs and static fusion methods still have clear limitations.

2.3 Research on the Application of Multimodal Fusion and Gating Mechanisms in Stock Price Forecasting

A single price sequence cannot fully reflect the different drivers of stock price volatility. For this reason, multimodal forecasting methods have begun to integrate price, text, and macroeconomic variables. Multi-source heterogeneous stock prediction research shows that multimodal fusion can improve the representation of market information^[4]. In general multimodal learning, the Gated Multimodal Unit learns how different modalities should influence a shared representation through multiplicative gates^[8]. In stock prediction, recent gated cross-attention research also indicates that gated fusion can help stabilize multimodal information when text data are sparse or semantically conflicting^[9]. However, few methods address the changing importance of text across trading days in a lightweight way. In stock forecasting, the price modality is usually more stable. The text modality is more vulnerable to noise and missing data. Therefore, a daily-level gating mechanism is useful. It can adjust the contribution of text according to the information available on each day.

3. RESEARCH METHODOLOGY

3.1 Overall Framework

This study proposes a stock price forecasting framework based on LSTM and cross-modal adaptive gating. The framework contains four parts. They are data processing, sentiment factor construction, model design, and experimental evaluation. First, daily stock data and stock forum comments are collected. They are aligned at the stock-trading-day level to form a daily sample table. Second, the text data are cleaned, scored, and aggregated by trading day. This step produces daily sentiment factors that reflect changes in investor sentiment. The factors are then combined with post counts and interaction metrics. This forms an influence-weighted representation. Third, three LSTM-based model categories are compared. They are Price-only LSTM, Price + Sentiment LSTM, and Gating-LSTM. The Gating-LSTM introduces daily gating coefficients. These coefficients adjust the contribution of text-group features. In the experimental evaluation, two logistic-regression models are also added as auxiliary non-LSTM baselines. Therefore, the later result section reports five model names while the methodological comparison still focuses on three LSTM-based categories. Finally, comparative experiments are used to evaluate the role of sentiment information and the cross-modal adaptive mechanism.

3.2 Data Processing and Construction of Sentiment Factors

The input data in this study include two main parts. The first part is historical A-share price data. The research objects include Ping An Bank (000001), BYD (002594), and Kweichow Moutai (600519). Their daily trading data are used as price features. The main variables include opening price, closing price, high price, low price, and trading volume. The second part is stock forum comment data. It includes text content, publication time, and interaction information. To avoid forward-looking information leakage, all data are aligned at the stock-trading-day level. The features are then merged into one daily sample table. After that, the text is cleaned. Sentiment scores are calculated for individual comments by using a simple sentiment lexicon. These scores are aggregated by trading day. The resulting variables include the positive sentiment ratio, negative sentiment ratio, and average sentiment score. They are also combined with post counts to construct daily sentiment factors.

3.3 Model Design

To test the effect of text features and the gating mechanism, this study compares three LSTM-based model categories. The first model is the Price-only LSTM. It uses only price features and serves as the baseline. The second model is the Price + Sentiment LSTM. It concatenates price features and daily sentiment factors before feeding them into the LSTM. This model tests whether text information can improve prediction. The third model is the Gating-LSTM. Its main idea is to adjust text features before fusing them with price features. Specifically, the price representation at time step t is denoted as h_{p^t} , and the sentiment representation is denoted as h_{s^t} . At the t -th time step, a joint representation is first constructed as $H_{ct} = [h_{p^t}; h_{s^t}]$. A linear layer and a Sigmoid function then generate the gating coefficient $g_t = \sigma(W_g H_{ct} + b_g)$. The gate value is then used to scale the sentiment representation: $\tilde{h}_{st} = g_t \odot h_{s^t}$. The scaled sentiment representation is concatenated with the price representation to form the final fused representation $H_{ft} = [h_{p^t}; \tilde{h}_{st}]$. Finally, H_f is fed into the LSTM encoder. A fully connected layer and a Sigmoid function output the predicted probability $\hat{y}$ that the stock will rise on the next trading day. This sentiment modulation process is illustrated in Figure 1. The overall architecture of the proposed Gating-LSTM is summarized in Figure 2. Compared with direct concatenation, this method can adaptively adjust the influence of the text modality according to the information strength on different trading days.




Figure 1. Simplified gating mechanism for sentiment modulation.




Figure 2. Simplified architecture of the Gating-LSTM model.

3.4 Experimental Setup and Evaluation Metrics

In the experiments, the training set, validation set, and test set are divided in chronological order. Sequence samples are constructed with sliding windows. The 2026 aligned dataset is relatively short. Therefore, the LSTM model uses a shorter lookback window in this stage. This study treats price direction prediction as a binary classification task. Accuracy and F1 are used as the main evaluation metrics. The average gate value of the gating model is also recorded. This helps to check whether the cross-modal gating mechanism actually participates in the model computations.

4. EXPERIMENTAL RESULTS AND ANALYSIS

4.1 Comparison of Main Results

This study tests five reported model names for stock price forecasting. These models include two auxiliary logistic-regression baselines and three LSTM-based categories: Logistic Regression, Price + Sentiment Logistic Regression, Price-only LSTM, Price + Sentiment LSTM, and Gating-LSTM. The preliminary experiments show an important problem. When

price data and text data are not aligned, text features are hard to use effectively. The results of different models also show little difference. Therefore, the data are re-modelled with price and comment data aligned to 2026. The first batch of experiments uses only Ping An Bank (000001). As shown in Table 1, the dataset is small. The comment coverage is also limited. Because of this, the models do not show clear improvement over the direct concatenation model, even after gating is added.

Table 1. Results for Ping An Bank.

Model	Test Accuracy	Test F1	Mean Gate
Logistic Regression	0.2500	0.0000	-
Price + Sentiment Logistic Regression	0.2500	0.0000	-
Price-only LSTM	0.0000	0.0000	-
Price + Sentiment LSTM	0.6667	0.8000	-
Gating-LSTM	0.0000	0.0000	0.4696

In the extension experiments, the dataset is expanded. Ping An Bank (000001) is retained, and BYD (002594) and Kweichow Moutai (600519) are added. This expansion improves the later experimental results. As shown in Table 2, the results of All-Stocks Logistic Regression and All-Stocks Price + Sentiment Logistic Regression are the same. Both models reach a validation accuracy of 0.50 and a validation F1 score of 0.4706. They also reach a test accuracy of 0.50 and a test F1 score of 0.5263. This means that sentiment features do not bring extra improvement to the linear model on the current pooled three-stock sample.

Table 2. Comparison of Linear Models.

Model	Validation Accuracy	Validation F1	Test Accuracy	Test F1
All-Stocks Logistic Regression	0.5000	0.4706	0.5000	0.5263
All-Stocks Price + Sentiment Logistic Regression	0.5000	0.4706	0.5000	0.5263

Table 3 shows the results of the standard LSTM models. The All-Stocks Price-only LSTM and the All-Stocks Price + Sentiment LSTM obtain the same test results. Both models achieve an accuracy of 0.2941 and an F1 score of 0.4545. This suggests that simple concatenation is not enough in the current setting. Adding sentiment features before the LSTM does not improve predictive performance for the time being.

Table 3. Comparison with a standard LSTM.

Model	Test Accuracy	Test F1
All-Stocks Price-only LSTM	0.2941	0.4545
All-Stocks Price + Sentiment LSTM	0.2941	0.4545

Table 4 compares the direct concatenation model and the gated model. The test accuracy of the All-Stocks Gating-LSTM is 0.7647. This is much higher than the 0.2941 achieved by the All-Stocks Price + Sentiment LSTM. However, its F1 score is 0.3333. This is lower than the 0.4545 of the direct concatenation model. The mean gate value of the Gating-LSTM is 0.5958. This shows that the gating module contributes to the model computations. Overall, the current results are mixed. The linear model and the standard LSTM do not gain much from sentiment features. The Gating-LSTM improves accuracy, but it does not improve F1. This suggests that the gating mechanism changes the model's classification bias.

Table 4. Comparison of direct concatenation and gated models.

Model	Test Accuracy	Test F1	Mean Gate
All-Stocks Price + Sentiment LSTM	0.2941	0.4545	-
All-Stocks Gating-LSTM	0.7647	0.3333	0.5958

4.2 Discussion of Results

The three result tables show that different models use text sentiment features in different ways. In the linear model and the standard LSTM, sentiment features do not lead to clear changes in performance. This suggests that the current sample size is still limited. The coverage of sentiment types is also not sufficient. As a result, the gain from text information is not yet stable. In contrast, the accuracy of the Gating-LSTM increases clearly. This indicates that the gating mechanism may improve overall classification accuracy. However, the lower F1 score also shows a problem. The model may be less balanced when distinguishing positive and negative classes. Overall, the results suggest that the cross-modal gating mechanism has potential. Its effectiveness still needs to be tested with larger samples and more complete data.

5. CONCLUSIONS AND FUTURE WORK

This study uses three A-share stocks as research objects. They are Ping An Bank (000001), BYD (002594), and Kweichow Moutai (600519). The study compares the roles of price features, text sentiment features, and cross-modal gating in stock movement prediction. The results show that sentiment features do not significantly change the performance of linear models. In standard LSTM models, direct sentiment concatenation also does not improve test accuracy or F1. This suggests that the benefit of text sentiment is still unstable in the pooled three-stock sample. By contrast, the test accuracy of the Gating-LSTM increases to 0.7647. This is higher than the 0.2941 achieved by the Price + Sentiment LSTM. However, its F1 score is 0.3333. This is lower than the 0.4545 of the Price + Sentiment LSTM. The result shows that the gating mechanism improves overall accuracy, but it still lacks balance in class identification. The mean gate value of the Gating-LSTM is 0.5958. This indicates that the gating module is involved in the model computations. Overall, the study verifies the feasibility of cross-modal gating in a multi-stock scenario. It also shows that model performance can differ across evaluation metrics. Future work can expand the stock pool and the time horizon. Future work can also increase the number of comment-covered trading days and improve the construction of sentiment factors. These steps would help assess the practical value of text information and gating mechanisms more comprehensively.

REFERENCES

- [1] Ling A F, Peng W, Wang Q Q, et al. A comprehensive review of research progress on the application of NLP to financial problems[J]. *Systems Engineering—Theory & Practice*, 2024
- [2] Wu Yanxin, Li Hongbin, Hu Guanzhen. A study on the optimisation of NLP-based stock selection strategies[J]. 2024
- [3] Xu Y, Cohen S B. Stock Movement Prediction from Tweets and Historical Prices[C]//*Proceedings of ACL*. 2018
- [4] Chen Qidong, Xie Jiayi, Cao Guangxi. A Model for Predicting Extreme Volatility in Small-Scale Stock Markets by Integrating Multi-source Heterogeneous Data[J]. 2025
- [5] Fischer T, Krauss C. Deep learning with long short-term memory networks for financial market predictions[J]. *European Journal of Operational Research*, 2018
- [6] Wu You. Research on a Stock Prediction System Based on an Improved LSTM Framework[D]. 2026
- [7] Liagkouras K, Metaxiotis K. A Hybrid Long Short-Term Memory with a Sentiment Analysis System for Stock Market Forecasting[J]. *Electronics*, 2025
- [8] Arevalo J, Solorio T, Montes-y-Gómez M, et al. Gated Multimodal Units for Information Fusion[C]. 2017
- [9] Zong C, Wan J, Cascone L, et al. Stock movement prediction with multimodal stable fusion via gated cross-attention mechanism [J]. *Complex & Intelligent Systems*, 2025