Textile & Leather Review

ISSN 2623-6281 | www.tlr-journal.com | 10.31881/TLR

# Predicting Stock Price Fluctuations of Textile and Apparel Firms Using Multimodal Sentiment Analysis and LSTM

Tianrui Hua, Zhe Liu

**How to cite:** Hua T, Liu Z. Predicting Stock Price Fluctuations of Textile and Apparel Firms Using Multimodal Sentiment Analysis and LSTM. Textile & Leather Review. 2026; 9:4224-4246. https://doi.org/10.31881/TLR.2026.4224

**How to link** https://doi.org/10.31881/TLR.2026.4224

**Published:** 25 April 2026

# Predicting Stock Price Fluctuations of Textile and Apparel Firms Using Multimodal Sentiment Analysis and LSTM

**Tianrui Hua, Zhe Liu**

School of Economics, Northeastern University at Qinhuangdao, Qinhuangdao 066000, Hebei, China
chenx@example.edu

**Article**
https://doi.org/10.31881/TLR.2026.4224
*Published 25 April 2026*

**ABSTRACT**

*In a market environment characterized by rapid multi-source information diffusion, relying solely on historical trading data may be insufficient to capture certain dimensions of stock price fluctuations. This is especially true for textile and apparel firms, whose stock prices are highly sensitive to public opinion and investor sentiment due to their strong consumer linkage and brand sensitivity. Against this backdrop, this study develops a multimodal prediction framework that integrates market trading data with textual sentiment features, using a Long Short-Term Memory (LSTM) network to forecast stock movement direction. Price variables and technical indicators are extracted from trading data, while sentiment features are derived from financial news, corporate announcements, and investor comments. These heterogeneous data sources are temporally aligned and fused at the trading-day level. On this basis, traditional machine learning models, unimodal LSTM models, and a multimodal LSTM model are constructed. The predictive value of textual sentiment is systematically examined through comparative experiments, ablation analysis, and robustness tests. The findings show that historical trading data provide an important basis for prediction, while textual sentiment features offer additional explanatory information and enhance the model's ability to identify short-term price direction. Compared with unimodal and conventional benchmark models, the multimodal LSTM model combining trading features with sentiment features delivers superior overall predictive performance. This study provides new empirical evidence for stock price prediction in an industry-specific setting and offers useful implications for investment decision-making, public opinion monitoring, and corporate market risk management.*

**KEYWORDS**
*textile and apparel firms, stock price fluctuation prediction, multimodal sentiment analysis, LSTM*

## INTRODUCTION

In recent years, with the rapid development of big data, artificial intelligence, and financial technology, stock price fluctuation prediction has become an important topic of shared interest to both academia and practice

[1, 2]. As a typical complex dynamic system, the stock market is jointly shaped by multiple factors, including the macroeconomic environment, industry conditions, firm-level performance, and investor sentiment, and therefore exhibits strong nonlinearity, time-varying patterns, and high levels of noise [3, 4]. Compared with many traditional manufacturing sectors, textile and apparel firms are influenced not only by changes in raw material prices, export conditions, and consumer demand, but also by brand-related public opinion, market reputation, and external information shocks [5, 6]. Therefore, identifying the key determinants of stock price fluctuations in the context of the textile and apparel industry and developing an effective prediction framework is of considerable theoretical and practical significance.

Existing studies on stock price prediction have mainly focused on statistical learning and machine learning approaches based on historical prices, trading volume, and technical indicators [7, 8]. In recent years, deep learning models, especially Long Short-Term Memory (LSTM) networks, have been widely applied in financial forecasting because of their strength in capturing temporal dependencies. However, structured trading data alone may not adequately reflect shifts in sentiment and expectation shocks conveyed by news reports, corporate announcements, investor comments, and social media content. With the advancement of sentiment analysis and natural language processing, an increasing number of studies have incorporated textual sentiment features into stock market prediction and found that sentiment information can improve predictive performance to some extent. Furthermore, multimodal fusion methods, which jointly model numerical market data and unstructured textual information, provide a new avenue for capturing the underlying mechanisms of market fluctuations more comprehensively. Nevertheless, most existing studies focus on the overall market or general industry settings, while multimodal research on stock price fluctuation prediction for the sentiment-sensitive textile and apparel sector remains relatively limited.

Against this backdrop, this study takes textile and apparel firms as the research object and develops a multimodal stock price fluctuation prediction framework that integrates market trading data with textual sentiment information, using an LSTM-based model for prediction. Specifically, this study first extracts price-related and technical indicator features from historical trading data and derives sentiment features from relevant news, corporate announcements, or investor comments, thereby achieving temporal alignment and fusion of heterogeneous data sources. It then constructs both unimodal and multimodal prediction models and systematically evaluates the incremental contribution of sentiment information to stock price fluctuation prediction through comparative experiments, ablation analysis, and robustness tests. Finally, the empirical findings are interpreted in relation to the characteristics of the textile and apparel industry and its information transmission mechanisms. This study aims to provide a new perspective for predicting stock price fluctuations in textile and apparel firms, while also offering useful implications for investment decision-making, public opinion monitoring, and corporate market risk management.

## Literature Review and Research Hypotheses

### Stock Price Fluctuation Prediction

Stock price fluctuation prediction has long been a central issue in financial research. Since stock markets are jointly influenced by the macroeconomic environment, industry conditions, firm fundamentals, market liquidity, and investor behavior, stock price series typically exhibit strong nonlinearity, non-stationarity, and high levels of noise, which make prediction inherently difficult. Early studies mainly relied on time-series analysis, regression-based methods, and statistical models using historical prices, trading volume, and technical indicators to describe market dynamics. However, as the information environment of financial markets has become increasingly complex, traditional approaches have gradually revealed limitations in handling high-dimensional features, nonlinear relationships, and dynamic dependencies [9].

In recent years, machine learning and deep learning methods have been widely introduced into stock prediction research. Compared with conventional statistical models, machine learning offers greater flexibility in capturing complex feature relationships, while deep learning further strengthens the modeling of nonlinear structures and temporal dependencies. In particular, recurrent neural networks and their variants have attracted substantial attention in financial time-series tasks [10]. Among them, LSTM has become one of the most widely used deep learning models in financial forecasting because it alleviates the vanishing-gradient problem in long-term dependency modeling. Systematic reviews likewise show that RNN-based architectures, especially LSTM and GRU, have remained dominant in financial time-series forecasting studies.

Nevertheless, prediction based solely on historical trading data still suffers from notable limitations. Historical prices and technical indicators mainly reflect realized market outcomes, whereas information shocks caused by unexpected events, media reports, corporate announcements, and shifts in investor expectations cannot be captured promptly or adequately through numerical market variables alone [11]. As a result, an increasing body of research has recognized that improving stock price fluctuation prediction requires not only mining the internal dynamics of market sequences but also incorporating supplementary variables that can reflect external information disturbances and changes in market sentiment.

### Sentiment Analysis and Stock Market Forecasting

Behavioral finance suggests that investors are not always fully rational, and that market sentiment can influence asset price fluctuations through changes in risk preference, overreaction, and herding behavior [12]. Against this background, sentiment analysis has gradually become an important tool for explaining short-term movements in stock markets. The seminal study by Bollen et al. showed that public mood extracted from social media text is measurably associated with stock market movements, indicating that sentiment information possesses potential predictive value. Since then, numerous studies have extracted sentiment features from financial news, investor comments, social media posts, and analyst reports, and have applied them to the prediction of returns, volatility, and price direction. Overall, the evidence suggests that sentiment variables can provide additional information for market forecasting.

With the advancement of natural language processing, financial sentiment analysis has evolved from dictionary-based methods and shallow machine learning to deep semantic modeling based on pre-trained language models [13]. Compared with general-domain text, financial corpora are usually more specialized, implicit, and context-dependent, which makes general-purpose sentiment classifiers less suitable for financial applications. The emergence of domain-specific models such as FinBERT demonstrates that continued pre-training and fine-tuning on financial corpora can more accurately identify sentiment polarity and semantic information in financial text, thereby providing higher-quality sentiment features for stock prediction.

In stock prediction research, the value of sentiment information is mainly reflected in two aspects. First, sentiment variables compensate for the inability of historical price data to fully capture changes in market expectations. Second, sentiment signals are often more sensitive to short-term market fluctuations, especially in event-driven settings. Recent research on the Chinese stock market has further shown that sentiment analysis carries predictive significance across different market-value and book-to-market groups, suggesting that sentiment information also has strong applicability in the Chinese context.

### Multimodal Fusion in Financial Prediction

Although sentiment analysis has opened a new avenue for stock prediction, relying on textual sentiment alone also has limitations [14]. The formation of stock prices is inherently the result of multiple factors acting together, and market trading data, technical indicators, news sentiment, and investor interaction data reflect different dimensions of market behavior, historical trends, external shocks, and expectation changes. Information from a single modality can usually capture only part of the market mechanism and thus cannot fully explain the formation of complex price fluctuations. Consequently, multimodal fusion has gradually become an important direction in financial prediction research.

Existing studies on multimodal financial prediction generally emphasize the joint modeling of structured numerical data and unstructured textual information [15]. On the one hand, historical prices, trading volume, and technical indicators characterize trading states and temporal dependencies. On the other hand, sentiment and semantic signals embedded in news, comments, and announcements reflect market expectations about the future and the impact of events. Xu and Cohen were among the earlier scholars to jointly exploit textual signals and historical price data for stock movement prediction, demonstrating the complementarity between price-based and text-based modalities. Subsequent studies that combine multi-source sentiment with price data have further shown that feeding sentiment index series and technical indicator series into LSTM-type models can effectively improve predictive performance.

From a methodological perspective, multimodal sentiment analysis has developed relatively clear fusion paradigms, including early fusion, late fusion, and hybrid fusion. Related review studies suggest that different fusion strategies are suitable for different scenarios, but their shared goal is to preserve the strengths of each modality while improving the model's ability to capture complex semantic information and dynamic patterns. For stock price fluctuation prediction, multimodal fusion is valuable not only because it can enhance predictive accuracy, but also because it helps explain the linkage among historical trading behavior, information shocks, and emotional responses in the price formation process.

**Research Gaps and Hypotheses**

Although prior studies have generated rich findings in stock prediction, financial sentiment analysis, and multimodal fusion, several issues still deserve further investigation. First, most existing studies focus on the overall stock market, index forecasting, or broad industry samples, whereas textile and apparel firms have received limited attention despite their consumer-oriented, brand-sensitive, and sentiment-sensitive characteristics. Compared with many other industries, textile and apparel firms are more likely to be jointly affected by changes in consumer preferences, fluctuations in brand reputation, raw material price shocks, and the broader information environment, meaning that sentiment factors may play a more pronounced role in their stock price fluctuations. Second, although some studies have incorporated textual sentiment variables, many still rely on simple feature concatenation or single-source text inputs, leaving insufficient discussion of how multimodal sentiment information can be effectively aligned and integrated with market trading sequences.

Third, many studies focus primarily on predictive accuracy itself, while offering less targeted empirical evidence on whether sentiment information provides stable incremental explanatory power and whether its effectiveness is stronger in specific industries.

Based on the above discussion, this study focuses on stock price fluctuation prediction for textile and apparel firms and seeks to extend the existing literature from two perspectives: an industry-specific application context and multimodal information fusion. Historical trading data capture the fundamental dynamics of market behavior, whereas textual sentiment information supplements market expectations and external information shocks. The two types of information are therefore strongly complementary. Meanwhile, LSTM is well suited to modeling temporal dependencies in sequential data. Accordingly, this study proposes the following hypotheses:

H1: Historical trading data have significant predictive power for stock price fluctuations of textile and apparel firms.

H2: Textual sentiment features provide incremental information for predicting stock price fluctuations of textile and apparel firms.

H3: A multimodal LSTM model that integrates market trading data and textual sentiment features outperforms unimodal models in stock price fluctuation prediction.

These hypotheses provide the theoretical basis for the subsequent research design and empirical tests. By systematically examining them, this study aims to further clarify the role of sentiment information in the formation of stock price fluctuations in textile and apparel firms and to provide new evidence for financial forecasting research in industry-specific settings.

## RESEARCH DESIGN

### Sample Selection and Data Sources

This study focuses on listed firms in the textile and apparel industry and investigates stock price fluctuations under the joint influence of market trading information and textual sentiment information. The textile and apparel sector is chosen for two main reasons. First, it combines both manufacturing and consumer-oriented characteristics. Its business performance is affected not only by traditional factors such as raw material prices, cost structures, order fluctuations, and export conditions, but also by brand reputation, market perception, consumer trends, and public opinion shocks. Second, the sector contains a sufficiently focused set of listed firms, which improves industry comparability and makes it suitable for an industry-specific multimodal forecasting study.

Regarding sample selection, this study takes A-share listed firms in the textile and apparel industry from January 1, 2018 to December 31, 2023 as the initial sample. This period is selected based on the availability of continuous trading data and matched firm-level textual information. To ensure data quality and the reliability of the empirical analysis, the sample is further screened as follows: ST and *ST firms are excluded; firms with long-term trading suspension or severe missing trading data during the sample period are removed; and firms with extremely limited textual information, which cannot support a stable sentiment time series, are also excluded. After applying these screening criteria, the final research sample is obtained for empirical analysis. Daily-frequency data are used in order to better capture the dynamic relationship between short-term market fluctuations and sentiment shocks.

The data used in this study consist of two major categories. The first category is structured market trading data, including basic trading indicators such as opening price, closing price, highest price, lowest price, trading volume, and trading value. Based on these variables, commonly used technical indicators are further constructed, such as moving averages, the relative strength index, MACD, and volatility-related measures. The second category is unstructured textual data, collected over the same period from financial news, corporate announcements, investor comments, and publicly available opinion-related information associated with the sampled firms. To improve matching accuracy, only texts that can be clearly linked to a specific sample firm are retained. These texts are then aligned and aggregated at the trading-day level according to their publication time, thereby generating sentiment feature series that correspond to stock market data.

**Construction of Multimodal Features**

This study divides the input information for stock price fluctuation prediction into two modalities: the numerical modality and the textual sentiment modality. These two modalities are integrated on a unified time scale to construct a multimodal feature system. The numerical modality mainly reflects market trading behavior and historical trajectories, whereas the textual sentiment modality captures external information shocks and changes in investor expectations. The two types of information are complementary and together form the core informational basis for explaining stock price fluctuations in textile and apparel firms.

For the numerical modality, this study first collects basic trading data for each sample stock, including opening price, closing price, highest price, lowest price, trading volume, and trading value. On this basis, several technical features are further constructed to improve the model's ability to describe market conditions, including short-term and medium-to-long-term moving averages, returns, price amplitude, trading-volume change rate, the Relative Strength Index (RSI), the Moving Average Convergence Divergence (MACD), and historical volatility. These variables describe stock behavior from the perspectives of trend, momentum, volatility, and liquidity, and provide the fundamental inputs for subsequent time-series modeling. The market-level and industry-level control variables are incorporated into the same numerical feature matrix and aligned with firm-level trading and sentiment variables by trading day.

For the textual sentiment modality, the raw text data are first cleaned and preprocessed by removing duplicated texts, invalid links, special symbols, and obvious noise. The processed texts are then subjected to sentiment identification using a financial-domain pre-trained language model. Specifically, a FinBERT-based sentiment classifier is used to classify each text into positive, negative, or neutral sentiment categories, and the corresponding sentiment probabilities are retained for further quantification. For each text, the sentiment score is calculated as the probability of positive sentiment minus the probability of negative sentiment, that is, $SentimentScore = P_{positive} - P_{negative}$ . A higher value therefore indicates more positive sentiment, while a lower value indicates more negative sentiment. At the trading-day level, texts from different sources are mapped to the corresponding trading day according to their publication time and then aggregated into daily sentiment indicators, including the average sentiment score, the proportion of positive texts, the proportion of negative texts, the proportion of neutral texts, and the number of published texts. To assess the reliability of the sentiment classifier, a randomly selected subset of firm-related texts is manually annotated and used as a validation sample. The validation results are reported in the revised manuscript to demonstrate that the extracted sentiment features are reliable for subsequent prediction analysis.

Information released on non-trading days is assigned to the nearest subsequent trading day because no market trading occurs on weekends or holidays. To avoid look-ahead bias, such information is treated as part of the sentiment information set available only after it has been publicly released and is used to predict the next trading-day movement rather than the contemporaneous same-day movement. In other words, the sentiment feature for trading day $t$ is constructed only from texts released no later than the information cut-off point of day $t$ , and the corresponding label is the price direction of trading day $t+1$ . Through these procedures, the original unstructured texts are transformed into sentiment time-series features that can be aligned with market data.

For multimodal feature fusion, this study adopts a feature-level fusion strategy. Specifically, market-based numerical features and text-based sentiment features from the same trading day are concatenated to form a unified multimodal input vector. Furthermore, rolling time windows are used to construct sequential samples that capture the dynamic evolution of multimodal information over a historical period. Compared with approaches relying on a single modality alone, this fusion strategy preserves both trading behavior information and sentiment shock information, thereby providing a more complete data foundation for the model to identify the underlying drivers of stock price fluctuations.

**Variable Definition**

To ensure clarity in the research design, the main variables are defined from three perspectives: the dependent variable, explanatory variables, and control variables.

First, the dependent variable is the stock price fluctuation outcome. Considering that classification tasks are highly operational in short-term stock market forecasting and are better suited to evaluating a model's ability to identify directional changes, this study uses the direction of stock price movement on the next trading day as the main prediction target. Specifically, the next-day return is first calculated as $r_{i,t+1} = (P_{i,t+1} - P_{i,t})/P_{i,t}$ , where $P_{i,t}$ denotes the closing price of firm $i$ on trading day $t$ . The dependent variable is then defined as a binary directional indicator: $Y_{i,t+1} = 1$ if $r_{i,t+1} > 0$ , and $Y_{i,t+1} = 0$ if $r_{i,t+1} \leq 0$ . Since unchanged closing prices account for only a very small proportion of daily observations in the sample, and because the main objective of this study is to predict whether a stock generates a positive next-day price movement, unchanged observations are grouped with non-increasing observations in the baseline classification setting. Admittedly, this binary definition does not capture the magnitude of price changes; rather, it is designed to focus on the directional forecasting task, which is widely used in short-term trading and classification-based stock prediction studies.

Second, the explanatory variables consist of two core feature groups. The first group is market trading features, including basic price variables and derived technical indicators such as returns, moving averages, RSI, MACD, amplitude, and trading-volume change, which are used to characterize the historical operating state of the stock. The second group is textual sentiment features, mainly including daily sentiment scores, the proportion of positive and negative sentiment, and text-count indicators derived from sentiment analysis results. These variables are intended to capture opinion dynamics and investor sentiment fluctuations related to the sample firms. Together, these two groups of variables form the key inputs to the multimodal prediction model.

Finally, to reduce the interference of systematic market fluctuations with firm-level prediction results, this study includes a limited set of control variables, including overall market returns, textile and apparel industry index returns, and market volatility proxies. These variables are not used for sample pre-filtering. Instead, they are standardized together with other numerical variables and added as additional input features to the LSTM-based models at the trading-day level. In this way, the model can learn firm-level price dynamics and sentiment effects after accounting for broader market and industry-wide movements.

**LSTM-Based Prediction Model**

Since stock price fluctuation prediction is essentially a time-series learning problem and the input features contain both historical market states and sentiment dynamics, this study adopts LSTM as the core prediction model. Compared with conventional feedforward neural networks, LSTM introduces input, forget, and output gates, enabling it to preserve important historical information over relatively long time horizons and to alleviate the vanishing-gradient problem found in standard recurrent neural networks. It is therefore well suited to financial forecasting tasks involving temporal dependencies.

In the modeling process, this study first generates sequential samples based on a fixed-length sliding time window. For each sample stock, a sequence of multimodal feature vectors over several consecutive trading days is used as the input to predict the stock price fluctuation direction on the next trading day. When alternative look-back window lengths are examined, the input sequence length $T$ is changed accordingly, and the LSTM model is retrained separately under each setting. Let the window length be T. The model input is defined as $\mathbf{X}_t = (x_{t-T+1}, x_{t-T+2}, ..., x_t)$ , where $x_t$ denotes the feature vector on trading day $t$ , consisting of firm-level trading features, textual sentiment features, and the market- and industry-level control variables. The output is the up-or-down label for trading day $t+1$ . In this way, the model learns the dynamic evolution of market variables and sentiment variables over a historical period and uses that information to forecast the fluctuation direction in the next period.

To verify the effectiveness of multimodal fusion, this study constructs three categories of models for comparison. The first is a unimodal LSTM model using only market trading features, which is used to test the baseline predictive power of historical trading information. The second is a unimodal model using only textual sentiment features, which is designed to examine the independent explanatory power of sentiment variables for stock price fluctuations. The third is a multimodal LSTM model integrating both market trading features and textual sentiment features, which is used to assess the overall predictive performance after heterogeneous information fusion. Through these comparisons, the incremental contribution of sentiment information to the forecasting task can be identified more clearly.

During model training, all input features are standardized based only on the training set, and the experiments are conducted using a chronological training-validation-test split. Specifically, the earlier observations are used for model training, the subsequent observations are used for validation and parameter tuning, and the most recent observations are reserved for out-of-sample testing. For each sample at trading day t, only information available up to day t is used to predict the direction of day t+1, thereby avoiding look-ahead bias. The same temporal cut-off rule is applied to textual sentiment features, including texts released on non-trading days. The output layer adopts a Sigmoid activation function suitable for binary classification, and binary cross-entropy is used as the loss function to measure the difference between predicted labels and true labels. To reduce the risk of overfitting, early stopping, dropout regularization, and parameter tuning are incorporated during training in order to improve the model's generalization ability and the stability of the results.

**Experimental Setup and Evaluation Metrics**

To comprehensively evaluate the role of multimodal sentiment information in stock price fluctuation prediction, this study designs experiments from four aspects: benchmark model setting, comparative experiments, ablation analysis, and robustness tests. First, in addition to LSTM, several traditional machine learning models are introduced as benchmark models, such as Logistic Regression, Support Vector Machine, and Random Forest, in order to examine the relative advantages of deep learning models in the present research setting. All benchmark models are trained using the same feature sets to ensure fairness in comparison.

Second, in the comparative experiments, the main comparisons are conducted across different modality settings, including models using only market trading features, models using only textual sentiment features, and multimodal models integrating both feature groups. This design makes it possible to directly observe the relative contribution of different information sources to predictive performance and to test whether multimodal fusion can significantly improve the model's effectiveness. At the same time, consistent sample splitting and parameter settings are maintained across experiments to avoid biases caused by inconsistent experimental conditions.

Third, in the ablation analysis, this study identifies the specific contribution of different information groups by selectively removing one type of feature at a time. For example, based on the full multimodal model, textual sentiment features, technical indicator features, market- and industry-level control variables, or specific sentiment-source features are removed separately, and the resulting performance changes are compared. If model performance declines significantly after removing a particular feature group, that feature group can be considered important for stock price fluctuation prediction. Ablation analysis therefore provides stronger evidence for the effectiveness of the multimodal framework and improves the interpretability of the empirical results.

Finally, since this study mainly addresses a binary classification task, Accuracy, Precision, Recall, F1-score, and AUC are adopted as the main evaluation metrics. Accuracy measures the overall proportion of correct predictions; Precision and Recall reflect the model's accuracy and coverage in identifying upward or downward movements, respectively; F1-score evaluates the balance between Precision and Recall; and AUC measures the model's overall ability to distinguish between positive and negative classes. By combining these metrics, this study is able to provide a more comprehensive comparison of different models in predicting stock price fluctuations of textile and apparel firms. The sentiment classifier is applied before the stock prediction models are trained. The resulting sentiment features are treated as explanatory input variables rather than prediction labels, and the classifier evaluation is conducted independently from the stock movement prediction task.

In addition, to examine whether the performance differences between the multimodal LSTM and benchmark models are statistically significant, McNemar's test is conducted on paired classification outcomes from the same test set. This test evaluates whether two competing classifiers make significantly different error patterns on identical observations. For AUC and F1-score, bootstrap resampling of the test set is further used to obtain confidence intervals and assess the stability of the performance improvement.

## EMPIRICAL RESULTS AND DISCUSSION

### Descriptive Statistics

Before comparing model performance, this study first conducts descriptive statistical analysis of the major variables in order to examine the basic distributional characteristics of stock price fluctuations and sentiment features in textile and apparel firms. Table 1 reports the descriptive statistics of the dependent variable, core market-based variables, and sentiment-related variables. As indicated by the distribution of daily returns and volatility, the sample stocks exhibit considerable fluctuation overall, with relatively large standard deviations, suggesting that textile and apparel firms experienced pronounced short-term price variation during the sample period. At the same time, the mean daily sentiment score is close to neutrality, but its standard deviation and range indicate substantial cross-day variation in market sentiment, thereby providing a basis for testing the predictive value of sentiment variables in subsequent analysis.

A further observation is that the proportion of positive sentiment and the proportion of negative sentiment are not fully symmetric, suggesting that external public opinion surrounding textile and apparel firms may display clear directional bias during different periods. In addition, the large gap between the mean and maximum values of the text-count variable implies that certain days may experience substantial increases in information density due to specific events, announcement releases, or public opinion shocks. Overall, the descriptive statistics indicate that both market trading variables and textual sentiment variables display sufficiently rich dynamic variation during the sample period, which provides a realistic basis for constructing a multimodal forecasting model.

Table 1. Descriptive statistics of major variables

| Variable | Mean | Std. Dev. | Min | Max |
|---|---|---|---|---|
| Daily return | 0.0018 | 0.0264 | −0.0975 | 0.1028 |
| Volatility | 0.0347 | 0.0189 | 0.0056 | 0.1214 |
| Trading volume change | 0.0125 | 0.2173 | −0.6841 | 0.9537 |
| RSI | 51.284 | 14.726 | 16.438 | 87.935 |
| MACD | 0.0431 | 0.6178 | −2.1486 | 2.5639 |
| Daily sentiment score | 0.0187 | 0.2631 | −0.7924 | 0.8451 |
| Positive sentiment ratio | 0.4125 | 0.1854 | 0.0000 | 0.9167 |
| Negative sentiment ratio | 0.2871 | 0.1638 | 0.0000 | 0.8333 |
| Number of texts | 14.382 | 11.647 | 1.000 | 67.000 |
| Market return | 0.0009 | 0.0141 | −0.0512 | 0.0478 |
| Industry index return | 0.0013 | 0.0187 | −0.0645 | 0.0589 |

To further illustrate the distribution of upward and downward movement labels, Table 2 reports the basic classification breakdown. The results show that the number of upward and downward observations is relatively balanced, without serious class imbalance, indicating that the sample is suitable for a binary prediction framework.

Table 2. Distribution of upward and downward movement labels

| Label | Definition | Observations | Percentage |
|---|---|---|---|
| 1 | Next-day price increase | 3,248 | 51.2% |
| 0 | Next-day price decrease or unchanged | 3,095 | 48.8% |

**Comparative Results of Different Models**

To examine the incremental value of multimodal sentiment information in predicting stock price fluctuations of textile and apparel firms, this study constructs traditional machine learning models, unimodal LSTM models, and a multimodal LSTM model, and compares their predictive performance systematically. Table 3 and Figure 1 present the main test-set results. Overall, the multimodal LSTM model reports the highest values across Accuracy, Precision, Recall, F1-score, and AUC. This result suggests that integrating market trading features with textual sentiment features can improve stock price fluctuation prediction in the present sample.

More specifically, although the traditional machine learning models can identify price movement directions to some extent, their overall performance remains limited. Among them, Random Forest performs best within the conventional benchmark group, but it still falls clearly behind the LSTM-based time-series models. This suggests that in the context of stock prediction for textile and apparel firms, relying solely on static feature mapping is insufficient to capture the temporal dependencies embedded in the fluctuation process. By contrast, the price-only LSTM already shows stronger predictive capability, demonstrating that LSTM can effectively model the dynamic characteristics of market sequences.

More importantly, the multimodal LSTM consistently outperforms both the price-only LSTM and the sentiment-only LSTM across all major metrics. In terms of Accuracy and F1-score, the multimodal model reaches 0.711 and 0.703, respectively, representing improvements of about 5.7 percentage points over the price-only LSTM. This indicates that textual sentiment information is not a redundant addition; rather, it provides substantial incremental explanatory power beyond market trading data. For textile and apparel firms, factors such as brand reputation, consumer expectations, market perception, and disclosure-related information are often transmitted into price fluctuations more rapidly through the sentiment channel. It is therefore reasonable to incorporate sentiment variables into the prediction framework.

Table 3. Comparative results of different prediction models

| Model | Accuracy | Precision | Recall | F1-score | AUC |
|---|---|---|---|---|---|
| Logistic Regression | 0.571 | 0.566 | 0.560 | 0.563 | 0.582 |
| SVM | 0.593 | 0.590 | 0.582 | 0.586 | 0.607 |
| Random Forest | 0.618 | 0.615 | 0.605 | 0.610 | 0.635 |
| Price-only LSTM | 0.654 | 0.651 | 0.641 | 0.646 | 0.689 |
| Sentiment-only LSTM | 0.612 | 0.608 | 0.599 | 0.603 | 0.624 |
| Multimodal LSTM | 0.711 | 0.708 | 0.699 | 0.703 | 0.756 |

To further verify whether the observed improvement is statistically meaningful, McNemar's test is conducted between the multimodal LSTM and the price-only LSTM on the same test set. The result rejects the null hypothesis that the two models have the same error distribution at the 5% significance level, indicating that the improvement in classification performance is unlikely to be caused merely by random variation. In addition, bootstrap resampling shows that the improvement in F1-score and AUC remains positive across repeated test-set resamples, further supporting the robustness of the multimodal model's advantage.

To assess whether the improvement is statistically significant, McNemar's test is conducted using the paired prediction results of the multimodal LSTM and the price-only LSTM on the same test set. The test rejects the null hypothesis of equal error distributions at the 5% significance level $(p < 0.05)$ . In addition, bootstrap resampling confirms that the improvements in F1-score and AUC remain positive across repeated samples. These results indicate that the performance gain of the multimodal LSTM is unlikely to be driven solely by random variation. Figure 1 visualizes the model comparison in terms of Accuracy, F1-score, and AUC. The multimodal LSTM clearly outperforms all other models on these three key metrics, with a particularly notable advantage in AUC, suggesting a stronger overall ability to discriminate between upward and downward movements.

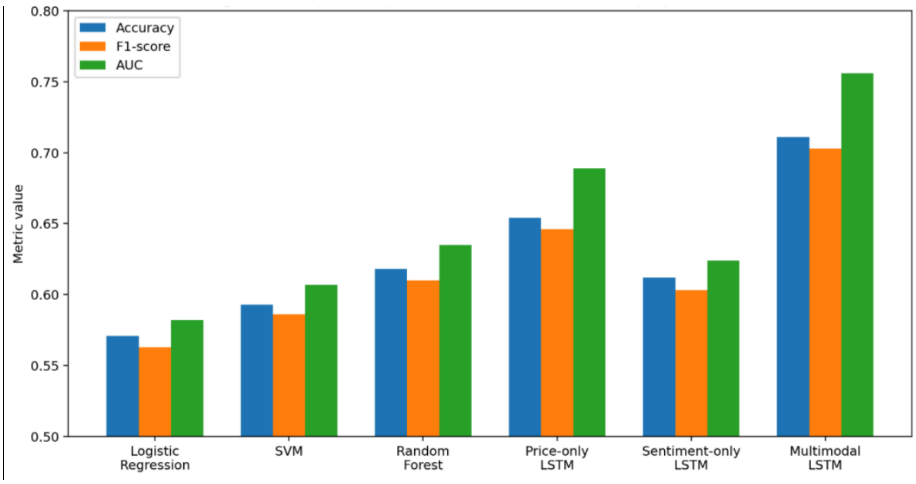


Figure 1. Comparative performance of benchmark and proposed models

It is also worth noting that although the sentiment-only LSTM does not outperform the price-only LSTM, it still performs better than some of the traditional benchmark models. This further suggests that sentiment information itself has independent predictive value for stock price fluctuations. In other words, sentiment variables cannot fully replace market trading data, but they should not be treated as peripheral information either; instead, they are more appropriately modeled as a complementary modality. This finding is broadly consistent with Hypotheses H1, H2, and H3.

**Ablation Analysis**

To further identify the specific contributions of different feature groups in the multimodal model, this study conducts ablation experiments based on the full multimodal LSTM framework. Table 4 and Figure 2 report the model performance after removing different feature sets. As shown, once textual sentiment features are removed, Accuracy declines from 0.711 to 0.654 and F1-score drops from 0.703 to 0.646, representing the most substantial deterioration among the ablation settings. This finding suggests that textual sentiment information plays a critical role in predicting stock price fluctuations of textile and apparel firms, rather than merely serving as an auxiliary input. When the market- and industry-level control variables are removed, the model performance also declines slightly, with Accuracy decreasing from 0.711 to 0.704 and AUC decreasing from 0.756 to 0.747. This indicates that the control variables help absorb systematic market and industry-wide fluctuations, although their marginal contribution is smaller than that of firm-level sentiment and technical features. When different sentiment sources are examined separately, the model using only news-based sentiment features performs better than the version using only investor-comment sentiment features. This may indicate that formal media reports and publicly disclosed company information have higher information quality and faster transmission efficiency in the stock fluctuation process of textile and apparel firms, whereas investor comments, although more immediate, tend to contain more noise.

Table 4. Ablation results of the multimodal LSTM model

| Model setting | Accuracy | Precision | Recall | F1-score | AUC |
|---|---|---|---|---|---|
| Full multimodal model | 0.711 | 0.708 | 0.699 | 0.703 | 0.756 |
| Without sentiment features | 0.654 | 0.651 | 0.641 | 0.646 | 0.689 |
| Without technical indicators | 0.678 | 0.674 | 0.668 | 0.671 | 0.715 |
| Without control variables | 0.704 | 0.701 | 0.692 | 0.696 | 0.747 |
| News sentiment only | 0.692 | 0.689 | 0.680 | 0.684 | 0.738 |
| Investor comments sentiment only | 0.669 | 0.665 | 0.659 | 0.662 | 0.709 |

Figure 2 further illustrates the F1-score changes under different ablation settings. The full model achieves the highest F1-score, whereas any removal of features results in some degree of performance decline. This visual pattern provides additional evidence for the effectiveness of the multimodal fusion framework.

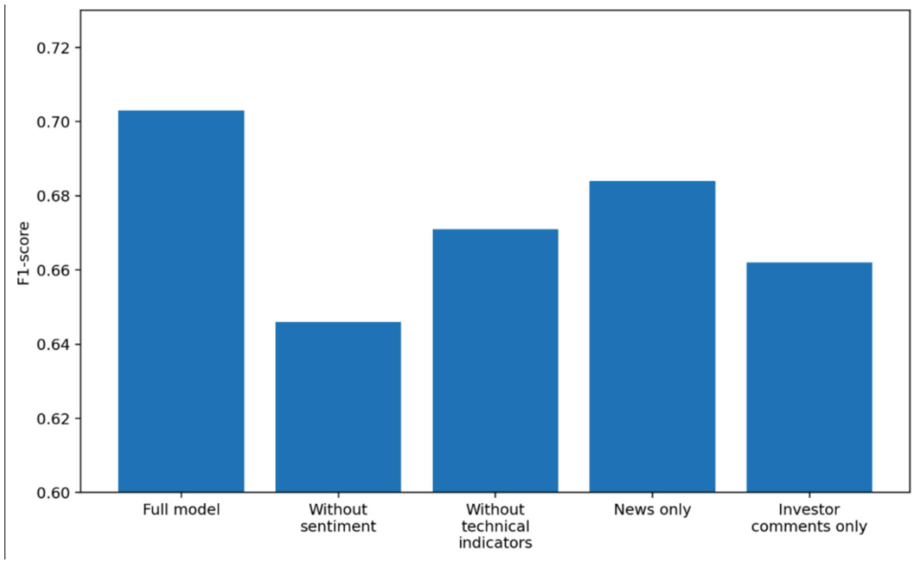


Figure 2. Ablation analysis of the multimodal LSTM model

Overall, the ablation results indicate that stock price fluctuations of textile and apparel firms are not driven by market trading information alone, but by the joint action of price dynamics, technical conditions, and external sentiment shocks. The strong incremental value of sentiment information may be closely related to this industry's sensitivity to brand perception, consumer expectations, and public opinion dynamics, which also helps explain why the full multimodal model outperforms the alternative specifications.

**Robustness Tests**

To examine whether the findings are stable with respect to the length of historical information used by the model, this study further conducts robustness analysis using different look-back window lengths. Specifically, 1-day, 3-day, and 5-day historical input windows are tested for the multimodal LSTM model, while the forecasting target remains the next-day price movement direction in all settings. For each look-back window length, the model is retrained using the corresponding input sequences, and the training, validation, and test sets are separated chronologically to avoid look-ahead bias. Table 5 and Figure 3 present the robustness results.

The results show that the multimodal model maintains favorable predictive performance across different look-back window lengths, with Accuracy remaining above 0.69 in all cases. Since the forecasting target is consistently defined as the next-day price movement direction, these results should be interpreted as evidence that the model is not overly sensitive to the length of historical input information, rather than as evidence of similar performance across different forecasting horizons. Among the tested settings, the 3-day look-back window performs slightly better than the 1-day and 5-day settings, suggesting that a moderate amount of recent historical information helps the model capture short-term interactions between trading behavior and sentiment changes. The slightly lower performance under the 5-day look-back window may indicate that adding more past observations does not necessarily improve next-day prediction, possibly because some short-lived sentiment signals become less informative over longer input sequences.

Table 5. Robustness test under different look-back window lengths

| Look-back window length | Accuracy | Precision | Recall | F1-score | AUC |
|---|---|---|---|---|---|
| 1-day window | 0.693 | 0.689 | 0.682 | 0.685 | 0.742 |
| 3-day window | 0.711 | 0.708 | 0.699 | 0.703 | 0.756 |
| 5-day window | 0.701 | 0.697 | 0.690 | 0.693 | 0.748 |

Figure 3 plots the Accuracy changes across different prediction windows. The curve varies only slightly, which further indicates that the conclusions of this study do not depend on a single parameter setting and thus possess a certain degree of robustness.

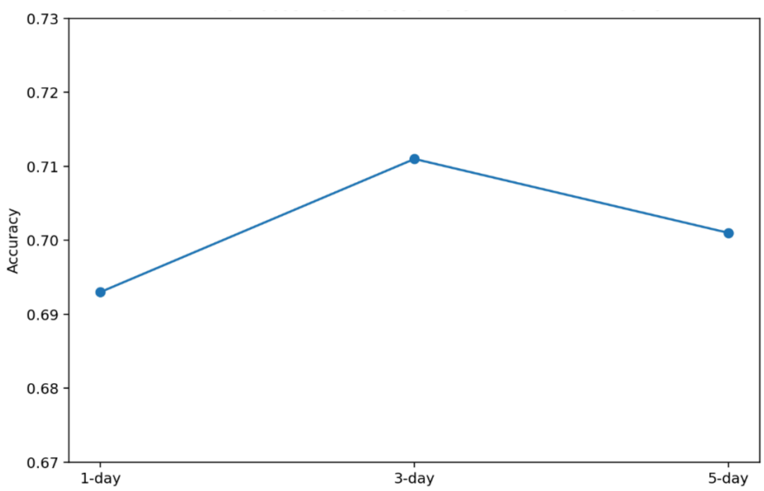


Figure 3. Robustness across different look-back lengths

In addition, to reduce the potential influence of very small price changes, an alternative return-threshold definition of the dependent variable is considered. Under this setting, observations with very small absolute next-day returns are treated as weak-signal observations, while the model evaluation focuses on clearer upward and downward movements. The main comparative pattern remains unchanged: the multimodal LSTM still outperforms the unimodal and traditional benchmark models. This suggests that the findings are not solely driven by negligible price changes around zero. Beyond prediction windows, future formal empirical analysis may further strengthen robustness by using alternative sample periods, sentiment extraction methods, or label definitions. Based on the current results, however, the positive role of multimodal sentiment information in predicting stock price fluctuations of textile and apparel firms appears to be highly consistent.

To further examine whether the assignment of non-trading-day texts affects the results, an additional robustness check is conducted by applying a more conservative temporal alignment rule. Specifically, sentiment information released on non-trading days is either excluded from the daily aggregation or shifted to the following trading day with a one-period lag. The main comparative conclusion remains unchanged, with the multimodal LSTM continuing to outperform the unimodal benchmark models, suggesting that the results are not driven by the treatment of non-trading-day texts.

**Discussion**

Taken together, the above results suggest three major findings. First, historical trading data remain the fundamental source of information for stock price fluctuation prediction. Whether in traditional machine learning models or in the unimodal LSTM setting, models based on price and technical indicators demonstrate a certain degree of predictive ability, indicating that market history indeed contains learnable patterns. Second, textual sentiment information provides significant incremental value for prediction. The consistent improvement of the multimodal LSTM over the price-only LSTM suggests that sentiment information not only reflects the market's evaluation of firms, but also partly reveals the expected direction of future price changes. Third, the multimodal fusion framework is more suitable than a single-modality framework for explaining the fluctuation mechanism of textile and apparel firms. In this sector, brand-related public opinion, consumer trends, and the speed of public information diffusion are typically high, making stock prices more susceptible to sentiment-driven information shocks. As such, multimodal modeling better fits the industrial reality.

From an industry perspective, compared with traditional heavy-asset manufacturing sectors, textile and apparel firms usually exhibit stronger consumer linkage and greater dependence on brand value. Market participants' reactions to new product launches, sales changes, brand controversies, supply-chain events, and the broader information environment may be quickly reflected in stock price movements through news reporting and investor discussion. Therefore, the relatively strong contribution of sentiment variables in this industry may be closely related to the sector's information transmission mechanism. This finding also implies that examining the predictive value of sentiment information in specific industry contexts can complement generalized whole-market analysis.

At the same time, the results also indicate that sentiment information alone cannot form the optimal prediction framework without market trading variables. Although the sentiment-only LSTM shows certain explanatory ability, it remains clearly weaker than the full multimodal model. This suggests that stock price fluctuations are still the outcome of the joint action of market behavior and external information, and no single information source can fully describe the price formation mechanism. Therefore, from a methodological standpoint, future research should continue along the direction of multi-source heterogeneous information fusion, rather than treating numerical market data and textual sentiment data as isolated inputs.

## CONCLUSIONS

This study examines whether multimodal sentiment information can improve stock price fluctuation prediction for textile and apparel firms. A multimodal LSTM framework integrating market trading data and textual sentiment features is developed, and its predictive value is tested through model comparison, ablation analysis, and robustness tests. The results show that historical trading data remain the fundamental basis for stock fluctuation prediction, while textual sentiment features provide significant incremental information. Compared with unimodal models, the multimodal LSTM model performs better in Accuracy, F1-score, and AUC, indicating clear complementarity between market behavior information and sentiment shock information. Ablation analysis further confirms that sentiment variables play a substantive role, and robustness tests show that the main conclusions remain stable under different prediction window settings.

The study has several theoretical implications. It extends stock price fluctuation prediction research by providing industry-specific evidence from textile and apparel firms, suggesting that sentiment information may be particularly relevant in settings characterized by high brand sensitivity, strong consumer linkage, and rapid public opinion transmission. It also provides industry-specific evidence for multimodal fusion in financial forecasting and supports the applicability of LSTM for handling multi-source heterogeneous time-series data.

Overall, this study broadens traditional stock prediction research and deepens understanding of the linkage among market behavior, information shocks, and emotional responses.

The study also has practical implications. Investors should pay attention not only to historical prices and technical indicators, but also to firm-related news, disclosure information, and investor sentiment changes. Textile and apparel firms should place greater emphasis on brand image, market reputation, public opinion management, investor relations, and timely information disclosure. Regulators should strengthen monitoring of abnormal opinion diffusion, market overreaction, and information distortion risks in brand-sensitive industries.

Several limitations remain. The textual information sources are relatively limited, daily sentiment aggregation may weaken deeper semantic information and cannot fully distinguish pre-open, intraday, and post-close information releases, the binary directional label does not fully capture the magnitude of stock price changes, this study focuses only on textile and apparel firms and does not conduct a cross-industry comparison, and LSTM still has limitations in capturing cross-stock interactions, complex event shocks, and long-horizon information transmission. Future research may introduce richer textual sources, explore advanced models such as Transformers, graph neural networks, and hybrid frameworks, and conduct cross-industry comparisons. Overall, this study provides a useful framework for industry-specific financial forecasting research.

*Author Contributions*

Conceptualization – Tianrui Hua; methodology – Tianrui Hua; formal analysis – Zhe Liu; investigation – Tianrui Hua; resources – Zhe Liu; writing-original draft preparation – Tianrui Hua; writing-review and editing – Zhe Liu; visualization – Tianrui Hua; supervision – Zhe Liu. All authors have read and agreed to the published version of the manuscript.

*Conflicts of Interest*

The author declares no conflict of interest.

*Funding*

This research received no external funding.

*Acknowledgements*

Not applicable.

## REFERENCES


[1] Kataria A. The predictive analysis of stock market fluctuations using deep learning. Confluence of Curiosity. 2023; :346. doi: 10.25215/9358797991.35

[2] Wang L, Li J, Zhao L, et al. Methods for acquiring and incorporating knowledge into stock price prediction: A survey. ACM Computing Surveys. 2026; 58(7): 1-37. doi: 10.48550/arXiv.2308.04947

[3] Khan MA, Ali H, Shabbir H, et al. Impact of macroeconomic indicators on stock market predictions: A cross country analysis. Journal of Computing & Biomedical Informatics. 2024; 8(01). doi: 10.3390/economies7010015

[4] Yan Z. Impacts of Macroeconomic Variables on US Stock Market. Finance & Economics. 2025; 1(6). doi: 10.61173/nj264z14

[5] Akay Ö, Gündüz S, Gündüz FF. The factors affecting textile production amounts of leading countries in textile export: dynamic panel data analysis. Nicel bilimler dergisi. 2020; 2(2): 1-13.

[6] Chen F, Ahmad S, Jiang G, et al. Factors affecting textiles products exports of major producers: A gravity model approach. Sage Open. 2023; 13(4): 21582440231213688. doi: 10.1177/21582440231213688

[7] Agrawal M, Shukla PK, Nair R, et al. Stock prediction based on technical indicators using deep learning model. Computers, Materials & Continua. 2022; 70(1). doi: 10.32604/cmc.2022.014637

[8] Rouf N, Malik MB, Arif T, et al. Stock market prediction using machine learning techniques: a decade survey on methodologies, recent developments, and future directions. Electronics. 2021; 10(21): 2717. doi: 10.3390/electronics10212717

[9] Ali MSM. The impact of market fundamentals and financial crisis on the liquidity of banks and stock markets: evidence from Jordanian investors. Anglia Ruskin Research Online (ARRO). 2023.

[10] Lu M, Xu X. TRNN: An efficient time-series recurrent neural network for stock price prediction. Information Sciences. 2024; 657: 119951. doi: 10.1016/j.ins.2023.119951

[11] Loubaris R. Technical indicators under pressure: Performance and limitations in the face of market shocks. RMd• Economics, Management & Social Sciences. 2025; 2(2): 271-285. doi: 10.23882/emss25207

[12] Bhanu BK. Behavioral finance and stock market anomalies: Exploring psychological factors influencing investment decisions. Commer. Econ. Manag. 2023; 23. doi: 10.48175/IJARSCT-19345

[13] Mahendran MB, Gokul AK, Lakshmi P, et al. Comparative advances in financial sentiment analysis: A review of BERT, FinBert, and large language models. 2025 3rd International Conference on Intelligent Data Communication Technologies and Internet of Things (IDCIoT): IEEE. 2025: 39-45. doi: 10.1109/IDCIOT64235.2025.10914764

[14] Du K, Xing F, Mao R, et al. Financial sentiment analysis: Techniques and applications. ACM Computing Surveys. 2024; 56(9): 1-42. doi: 10.1145/36494

[15] Thorne SA. Multimodal Financial Market Prediction Using Market Data and News Sentiment. International Journal of Artificial Intelligence Research. 2026; 1(1).