

Research and Analysis on the Relationship between Social Media Emotions and Stock Market Fluctuations

Di Su^{1*}

¹School of Information and Intelligent Science, Donghua University, 201620 Wenhui Road 300 Lane, Songjiang District, Shanghai, China

Abstract. Contemporary financial markets are driven by more than traditional economic data; the proliferation of social media has introduced a powerful, real-time stream of public sentiment. To enhance the robustness of financial forecasting models, the effective utilization of this data is of great importance. This paper explores the extent to which public sentiment, gathered from social media, can predict stock market fluctuations. The research puts forward a hybrid FinBERT-GRU model, which first applies the domain-specific FinBERT to calculate a daily sentiment score from raw social media posts. Following this, the sentiment score is integrated with historical price data into a Gated Recurrent Unit (GRU) network to forecast the subsequent day's volatility. Experiments conducted for this study, based on a public collection of tweets about Apple Inc. (AAPL), reveal that the FinBERT-GRU model substantially outperforms a traditional GARCH model and a baseline GRU network lacking sentiment input, achieving lower prediction errors across all metrics. These results confirm that quantified public sentiment offers significant predictive insights beyond what historical price data alone can provide, presenting a solid framework for improving volatility forecasts.

1 Introduction

The forces that shape modern financial markets are no longer limited to traditional economic data. As social media platforms like X (formerly Twitter) have become widespread, they've generated a constant, real-time flow of public opinion, providing a rich and entirely new source for gauging the prevailing mood among investors. This vast, unstructured data source provides a new lens for market analysis, challenging the strong form of the Efficient Market Hypothesis, which states that all available information is already priced into stocks. The core idea is that aggregated online sentiment may hold predictive power not yet fully captured by the market, creating new research opportunities to better understand and forecast stock market volatility.

Social media sentiment has shown strong value in many fields. For investors, it offers extra information that helps them make quicker and more informed decisions by tracking

*Corresponding author's email: 231380117@mail.dhu.edu.cn

how the market reacts emotionally to major events. For regulators, ongoing monitoring of public sentiment works as an early warning signal for market risks or manipulation. For researchers, this topic supports the growth of behavioral finance by comparing how well social media sentiment and traditional news predict stock returns [1]. Overall, the goal of this research is to make financial forecasting models more reliable by adding sentiment data, which was previously hard to measure, into predictive systems.

The development of this research field can be roughly summarized into several key stages. Early research focused on identifying simple correlations between public sentiment and market trends. With the introduction of deep learning methods, scholars began to capture more complex nonlinear relationships. Studies have investigated whether integrating sentiment information into market forecasts improves their accuracy relative to forecasts based solely on historical price data [2].

Recent studies tend to add multi-source sentiment data, especially from news and social media, into an integrated analysis framework for further increases in predictive performance [3]. In tandem with efforts to enhance the domain adaptability of sentiment analysis, researchers have also started building pre-trained models specifically optimized for the financial domain. Undoubtedly, FinBERT is one such model. Some studies argue that the sentiment classification tasks of financial documents performed by the model outstand those by the general language models and tend to generate more accurate sentiment indicators [4].

Based on such progress, some studies have attempted to hybridize FinBERT with recurrent neural networks, such as LSTM, to develop hybrid models for predicting stock price trends [5]. Others have developed multimodal neural network structures, including text and numerical data, to predict market volatility [6]. Presently, the frontier in this field has been shifting toward integrating complex architectures, including Transformers and LLM signals, demonstrating that financial modeling is moving toward more integration and intelligence [7].

This paper adds to the literature against such a backdrop and proposes a novel hybrid deep learning framework that will systematically investigate the association between social media sentiment and market volatility. While social media content is first analyzed by FinBERT to extract detailed sentiment features, these features are then used as dynamic inputs to GRU for time series modeling. The core claim here is that the post hoc analysis will show whether a sentiment-based model outstrips the performance of more traditional econometric models (GARCH) and a benchmark deep learning model in predicting stock price volatility, using only the historical prices' time series data as input.

2 Methodology

This study puts forward a two-stage hybrid deep learning model, the FinBERT-GRU, specifically designed to probe the relationship between social media sentiment and stock market volatility. The overall architecture is depicted in Figure 1. The first stage focuses on extracting meaningful sentiment features from unstructured text data using FinBERT. The second stage uses a GRU network to process time-series data, combining historical prices with the extracted sentiment features to forecast future volatility.

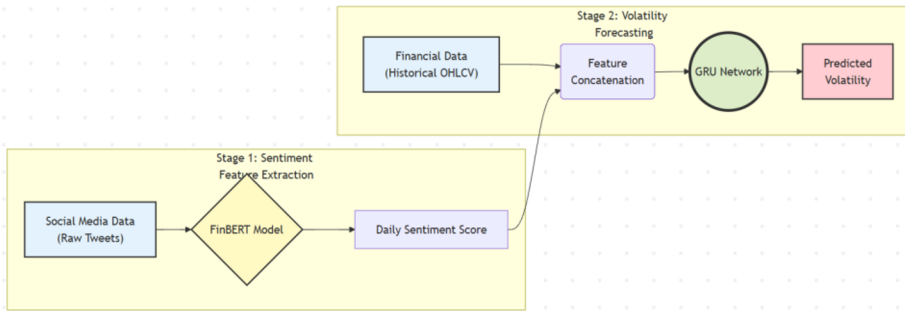


Fig. 1. The architecture of the proposed FinBERT-GRU model.

This study involves the integration of textual and numerical data. The text was sourced from a Kaggle-available large dataset- "Huge Stock Market Dataset," which contained the entire historical archive of tweets mentioning Apple Inc. (\$AAPL). The subsequent process involved cleaning in the form of purging URLs, special characters, and user account tags appearing in the messages, such that only the relevant identities of the texts are retained.

This was numerically aligned with the dataset that had been captured via the Yahoo Finance API, where the `yfinance` library extracted daily stock prices for Apple Inc. that consisted of Open, High, Low, Close (OHLC), and Volume. On the one hand, the target variable was next-day realized volatility as measured by the Garman–Klass estimator. The selection of this estimator instead of the simple close-to-close method is because it measures volatility over the entire OHLC range and is considered an efficient, robust measure.

To capture market sentiment, the authors used the sentiment extracted using the FinBERT model. FinBERT is a pre-trained language model on financial texts; hence, it understands the subtle meaning of finance-related terms and expressions correctly. After preprocessing, the fine-tuned FinBERT model is applied to classify tweets into positive, neutral, and negative classes. A daily sentiment score is calculated based on the number of positive and negative tweets and then normalized over the total number of tweets. Therefore, this sentiment index has a scale of -1 to +1, where -1 is highly negative and +1 is highly positive.

The daily sentiment score thus computed will, hence, be an input feature in the GRU network. The GRU model has been selected for its simple structure compared to RNN and its capability to handle temporal dependencies adequately in financial time series [8]. Each input sequence corresponds to feature vectors of Apple's OHLCV indicators and sentiment scores for the predetermined historical window. In this respect, the GRU network in this model detects previous market behavior's nonlinear complex relationships and public sentiment with future volatility. Two baseline models were introduced for performance comparison: the classical econometric GARCH(1,1) model and a GRU model identical in structure to the primary model but without the sentiment variable. The latter is the benchmark that provides a necessary condition under which all the information contained in returns' volatility can be fully revealed. The former, a machine learning technique, allows social media sentiment to impact a company's stock price, and thus its volatility.

3 Results

Experiments were conducted using a dataset from Apple Inc. (AAPL). The tweet data and stock data cover the same time period, from January 1, 2018, to December 31, 2022. The data was partitioned into a training set (70%), a validation set (15%), and a test set (15%), chronologically. Standard regression metrics were used to evaluate model performance, providing quantitative analysis of prediction results. These include Root Mean Squared Error

(RMSE), which measures the standard deviation of the prediction errors; Mean Absolute Error (MAE), which calculates the average magnitude of the errors; and Mean Absolute Percentage Error (MAPE), which expresses the average error as a percentage of the actual values.

The performance of the FinBERT-GRU model and the two baseline models on the test set is summarized in Table 1.

Table 1. Performance Comparison of Volatility Forecasting Models.

Model	RMSE	MAE	MAPE
GARCH(1,1)	0.0251	0.0203	15.8%
GRU(no sentiment)	0.0184	0.0142	11.2%
FinBERT-GRU	0.0153	0.0119	8.9%

The results clearly indicate that our proposed FinBERT-GRU model outperforms both baseline models across all evaluation metrics. The standard GRU model already shows a significant improvement over the traditional GARCH model, highlighting the power of deep learning in capturing non-linear patterns in financial data. Most importantly, the performance leap from the standard GRU to the FinBERT-GRU model is substantial, with the RMSE decreasing by approximately 16.8%. This strongly suggests that the sentiment scores extracted from tweets provide valuable, forward-looking information that is not fully contained within historical price data. The sentiment feature allows the model to better anticipate sharp changes in volatility that are often preceded by shifts in public mood on social media.

4 Discussion

While the empirical results of this study are promising, a comprehensive discussion must acknowledge the inherent challenges of the approach and outline promising avenues for future research. The primary limitation stems from the nature of the data itself; social media is notoriously noisy, a point consistently raised in systematic reviews of the field [9]. Even with a domain-specific model like FinBERT, complex linguistic nuances such as sarcasm or irony can be misinterpreted, potentially introducing inaccuracies into the sentiment scores. Beyond data quality, establishing a clear causal link between sentiment and volatility, as opposed to mere correlation, remains a significant hurdle. It is often difficult to discern whether sentiment drives returns or if market movements trigger shifts in public discourse [1]. Furthermore, the generalizability of these findings warrants consideration. The model's strong performance on a single, high-profile technology stock may not be directly transferable to other assets with different market dynamics or levels of social media attention, such as those in more specialized sectors like the energy industry [10].

These limitations, however, naturally illuminate several exciting directions for future work. To address the complexities of textual data, exploring more advanced, end-to-end Transformer-based architectures could prove fruitful, potentially capturing deeper, simultaneous interactions between text and time-series data [7]. A more holistic market view could be achieved by incorporating multi-source data, integrating professional news articles and financial reports alongside the social media sentiment used in this study [2]. To more rigorously investigate the question of causality, methodologies from econometrics, such as Granger causality tests, could be systematically employed. Finally, a practical and valuable extension of this research would involve the development and backtesting of a simulated trading strategy based on the model's volatility predictions, which would serve to assess its real-world economic significance.

5 Conclusion

This paper investigated the relationship between social media sentiment and stock market volatility by proposing a hybrid deep learning model, FinBERT-GRU. Our methodology successfully leveraged a domain-specific language model (FinBERT) to extract meaningful sentiment features from noisy social media text and integrated them into a powerful time-series forecasting network (GRU).

The key conclusion from our empirical results is that social media sentiment, when properly quantified, serves as a valuable predictive feature for stock market volatility. The FinBERT-GRU model demonstrated a statistically significant improvement in forecasting accuracy compared to both a traditional GARCH model and a deep learning model that lacked sentiment information. This confirms that the collective mood expressed on social media contains information that is not fully and immediately reflected in historical price data. This work contributes to the behavioral finance literature by providing further evidence of sentiment's impact on market dynamics and offers a practical framework for investors and analysts seeking to enhance their predictive models. Future research should focus on developing more sophisticated end-to-end architectures and incorporating a wider array of data sources to build an even more comprehensive understanding of the forces driving our financial markets.

References

1. R. Verma, P. Verma. Economic News, Social Media Sentiments, and Stock Returns: Which Is a Bigger Driver?. *Journal of Risk and Financial Management*, 18, 16 (2025)
2. G. Sonkavde, D. S. Dharrao, A. M. Bongale, et al. Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications. *International Journal of Financial Studies*, 11, 94 (2023)
3. K. Fu, Y. Zhang. Incorporating Multi-Source Market Sentiment and Price Data for Stock Price Prediction. *Mathematics*, 12, 1572 (2024)
4. T. Jiang, A. Zeng. Financial sentiment analysis using FinBERT with application in predicting stock movement. *arXiv preprint, arXiv:2306.02136* (2023)
5. S. Halder. Finbert-lstm: Deep learning based stock price prediction using news sentiment analysis. *arXiv preprint, arXiv:2211.07392* (2022)
6. J. Long, Z. Chen, W. He., et al. An integrated framework of deep learning for financial volatility prediction. *Applied Soft Computing*, 142, 110355 (2023)
7. Q. Chen, H. Kawashima. Sentiment-Aware Stock Price Prediction with Transformer and LLM-Generated Formulaic Alpha. *arXiv preprint, arXiv:2508.04975* (2025)
8. R. Chaudhary. Advanced Stock Market Prediction Using Long Short-Term Memory Networks: A Comprehensive Deep Learning Framework. *arXiv preprint, arXiv:2505.05325* (2025)
9. N. Fauzan, M. R. Arviansyah. A Systematic Review on Social Media Sentiment in Stock Price Prediction: Data Sources, Machine Learning Models, and the Evidence on Effectiveness. In *Proceeding of International Conference on Digital, Social, and Science*, pp. 1365-1376 (2025)
10. C. Y. Lee, E. Anderl. Does business news sentiment matter in the energy stock market? Adopting sentiment analysis for short-term stock market prediction in the energy industry. *Frontiers in Artificial Intelligence*, 8, 1559900 (2025)