

RESEARCH ARTICLE

Stock Price Prediction with LLM-Guided Market Movement Signals and Transformer Model

Qizhao Chen^{1,*} ¹Graduate School of Information Science, University of Hyogo, Japan

Abstract: Accurate stock price prediction is difficult because financial markets are complex and affected by various factors. Traditional methods often fail to analyze financial news and market trends effectively. Recent improvements in large language models (LLMs) allow better extraction of insights from text and stock data. Motivated by these advances, this paper introduces a new approach that combines LLM-generated stock predictions with a Transformer model to forecast prices. Specifically, a structured prompt fed with financial news and stock features is designed. Then, the prompt-based LLM approach is used to generate the predicted trend in the form of a one-hot vector and its corresponding probability. Finally, the LLM-generated output, together with historical closing prices, is used as input to a Transformer model to predict the stock price for the next day. To examine the effectiveness of the proposed framework, models such as Long Short-Term Memory (LSTM), Temporal Convolutional Network, Convolutional Neural Network (CNN), CNN-LSTM, Random Forest, support vector regressor, XGBoost, and vanilla Transformer are chosen for comparison. In addition, an ablation study is conducted under several configurations. The findings indicate that the proposed framework exhibits superior performance compared to all the baseline models. Moreover, the ablation study demonstrates that integrating LLM-predicted features has the potential to improve stock price prediction performance.

Keywords: deep learning, financial news, LLM, stock price prediction, sentiment analysis, transformer

1. Introduction

Stock price prediction is a popular but challenging research area in finance, computer science, and statistics because an accurate stock price prediction enables investors and traders to make a large amount of money from their investments, but the stock price is affected by many factors, such as market sentiment, historical trends as well as macroeconomic factors. Moreover, the news of one company can affect others as well, especially when they have close business ties. Therefore, capturing the accurate movement of stock prices is difficult. Although traditional models such as Long Short-Term Memory (LSTM), support vector machine (SVM), and Random Forest (RF) can roughly capture the stock movements using historical data, they often fail to capture the intricate relationship between stock prices and market movements due to the sudden change in market sentiment, which can significantly influence future price behavior.

Recently, the rapid development of natural language processing (NLP) and large language models (LLMs) has provided new opportunities to integrate unstructured textual data into stock price prediction models. Some important sources of this kind of unstructured textual data include financial news and social media posts, such as tweets and comments from some websites. News and social media posts are important sources to provide insights to reflect the

underlying market sentiment. For example, tweets from US President Donald Trump often led to immediate fluctuations in the stock market. These reactions illustrate how sentiment embedded in text-based sources can drive market behavior in real time. However, these signals are often overlooked in traditional stock price prediction models. The ability to process large amounts of text data allows LLMs to extract valuable information, such as sentiment polarity and specific market events, which can be used to complement numerical data to improve stock price prediction.

In this paper, we introduce a novel framework of combining LLM-generated features with a Transformer model for stock price prediction. The Transformer models, originally designed for NLP tasks, have shown good performance in capturing long-term relationships within sequential data. Unlike recurrent neural networks (RNNs) such as LSTM, which process data sequentially and may struggle to capture long-term relationships, the Transformer model uses self-attention mechanisms to process all input features simultaneously. This allows it to efficiently learn complex relationships in stock price movements while also incorporating external market sentiment.

Our proposed approach improves traditional stock price forecasting by considering LLM-generated sentiment features, including stock movement predictions (Up, Down, or Same) and their corresponding probabilities. These sentiment-driven insights serve as additional input features, enabling the Transformer model to consider both historical stock trends and market sentiment when making predictions.

*Corresponding author: Qizhao Chen, Graduate School of Information Science, University of Hyogo, Japan. Email: af24o008@guh.u-hyogo.ac.jp

This study aims to answer the following research question: can LLM-generated features improve the predictive accuracy of Transformer-based models for stock price forecasting?

The main contributions of this research are as follows:

- 1) We introduce a new approach that combines Transformer models with LLM-generated stock predictions, including their confidence levels, to better predict price movements. This helps the model understand market sentiment and trends more effectively while making more accurate forecasts.
- 2) We conduct an ablation study with different configurations to demonstrate the effectiveness of combining LLM-generated sentiment features with historical stock prices. The ablation study shows that the full model outperforms baselines such as LSTM, support vector regressor (SVR), and RF, achieving superior performance in terms of mean squared error (MSE).

The remainder of this paper is structured as follows. First, related work is listed (Section 2). The methodology is then described (Section 3). In Section 4, the experimental results are presented. Section 5 illustrates the ablation study. Finally, Section 6 concludes the paper.

2. Literature Review

Traditional machine learning models such as SVMs [1–3], RF [4–6], and gradient boosting [7–9] are popular candidates for stock price prediction. These models can better predict stock movements by identifying complex patterns in market data. For example, SVM has been widely applied in stock movement classification tasks, while RF and gradient boosting have been used for regression-based stock forecasting. However, these models often fail to capture long-term relationships within time series data.

With the emergence of deep learning, researchers have increasingly adopted LSTM [10–12] and Transformer-based architectures [13–15] for financial forecasting. LSTM models work well for analyzing stock price trends because they are specifically designed for sequential data processing. The main disadvantages are their slower processing speed and expensive computational costs.

The Transformer model, initially introduced and designed for NLP tasks, has recently been applied to financial time-series prediction. Unlike LSTMs, the Transformer processes input sequences in parallel using self-attention mechanisms, allowing it to capture both local and global dependencies effectively. Several studies have explored the use of vanilla Transformers and transformer-based architectures such as Informer models [4, 16, 17] for stock forecasting, demonstrating their superior performance compared to recurrent architectures. However, most Transformer-based stock prediction models focus solely on historical price movements and technical indicators, without incorporating external sentiment-driven signals.

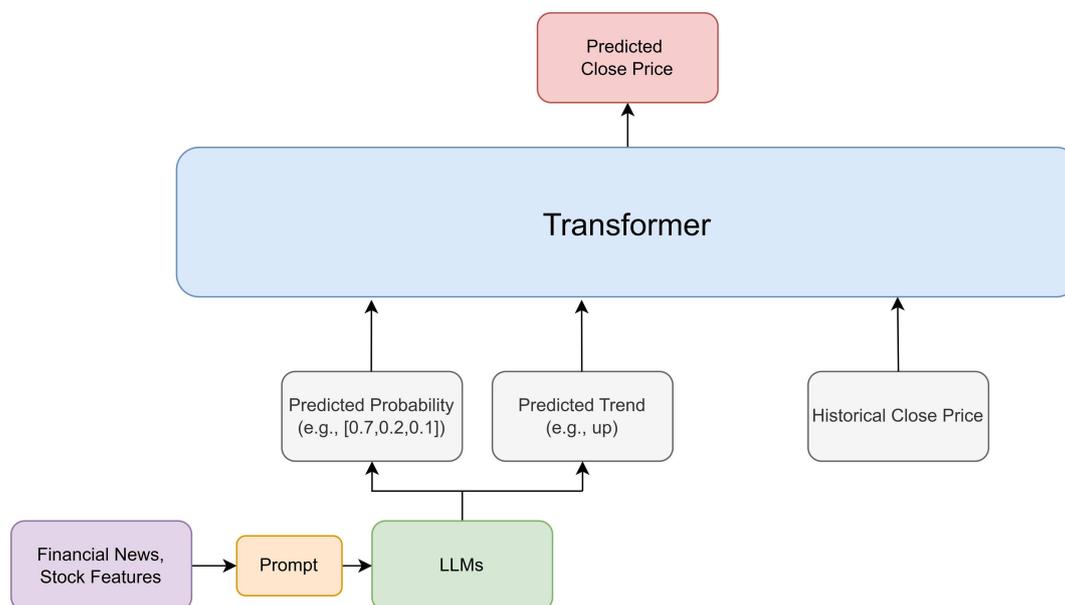
Recently, a growing number of research have explored the integration of LLM-generated sentiment features into stock prediction models [18]. Some studies apply fine-tuned BERT models to classify news sentiment [19, 20], while others use GPT-style models to extract market insights from financial documents [21, 22]. However, these approaches often require fine-tuning on domain-specific datasets, making them computationally expensive. In contrast, prompt-based LLM frameworks offer a more flexible solution by allowing models to generate structured outputs based on carefully designed textual prompts. This method eliminates the need for extensive fine-tuning while still leveraging the advanced reasoning capabilities of LLMs.

The latest trend is to utilize prompt-based LLMs directly to predict stock prices, such as FinMA [23] and FinGPT [24]. However, the direct application of prompt-based LLMs exhibits limitations. For example, Chen [25] shows that although LLMs can generate a relatively accurate prediction of stock price magnitude, they fail to capture the right stock movement (up or down) for half of the time. To improve the performance of LLMs in time series forecasting, some new methods are needed.

3. Research Methodology

Figure 1 shows the overall picture of the proposed framework. First, each prompt, which consists of both financial news and stock features with a sliding window size of five days, feeds

Figure 1
Proposed LLM-transformer framework



the LLM. Then, the LLM-generated output, including the predicted stock movement and its probability, together with the historical close prices with a sliding window size of five days, is used as input to the Transformer model. The Transformer model finally outputs the predicted close price for the next day.

3.1. Data

In this study, we collect financial data and news data related to the target stock as well as its corresponding related companies' stocks to construct a comprehensive dataset for prediction. Related companies here are defined as those that have close business ties or other relationships with the target companies such as suppliers and business partners. For example, for Apple Inc., news articles from companies such as Amazon, Google, Meta, and Samsung are also collected.

The stock market data include Open, High, Low, Close, and Volume (OHLCV), while financial news data provides additional contextual information. Daily stock data are downloaded from Yahoo Finance from 01/05/2015 to 08/05/2024 using the Python library `yfinance`. Yahoo Finance is selected as the data source due to its availability, ease of access, and comprehensive coverage of historical stock prices in different financial markets around the world. It is often used in academic and industry research, which ensures reproducibility and consistency. Daily news data of the corresponding companies from the same period are downloaded from some online websites such as Yahoo News using the Financial News Feed and Stock News Sentiment data API¹. The target stocks collected in this study include Apple (AAPL), HSBC (HSBA.L), Tencent (0700.HK), and Toyota (7205.T), which represent globally recognized companies from different major markets: the United States, Europe, China, and Japan. This diversity allows the model to be tested across different market environments, which can improve the generalizability of the findings. These companies are also chosen due to their large market capitalization, high trading volume, and significant media coverage, which ensures the availability of both financial data and related textual data, such as news articles.

To structure the dataset, we apply a sliding window approach, using the past five days of historical stock prices as sequential input for both the LLM and transformer. This allows both the LLM and transformer to learn temporal dependencies in stock price movements. The five-day sliding window can capture the short-term market dynamics within a typical trading week. The first 70% of the data will be used for training, and the remaining 30% of the data will be used for testing.

3.2. LLM for stock movement prediction

To integrate financial news and stock data into stock price prediction, we apply a prompt-based LLM approach to forecast the stock price. The LLM is fed with a structured prompt (Figure 2) containing financial news headlines and numerical stock features from the past five trading days. This prompt ensures that the model considers both historical price trends and qualitative sentiment when predicting stock movement. Specifically, in designing the prompt, the goal is to have the LLM emulate the reasoning process of financial analysts, who consider both quantitative market data and qualitative sentiment information. To achieve this, we structured the prompt to direct the model's attention across multiple dimensions, including company-specific news, related industry developments,

Figure 2
Sample prompt used in DeepSeek model

```
Analyze historical stock data and relevant news data to predict whether the closing price of Apple (AAPL) will rise or fall on {next_date}.
Consider the following factors:
  • News directly related to Apple (e.g., earnings reports, product launches, executive changes).
  • News about companies with strong business ties or market correlations to Apple, such as major suppliers, competitors, and industry partners.
  • Broader market trends and macroeconomic factors that may impact Apple's stock price.
Stock and news data for your reference: {stock and news dataset}
Output Format (JSON):
{
  "prediction": "Up" | "Down" | "Same",
  "probabilities": {
    "Up": 0.XX,
    "Down": 0.XX,
    "Same": 0.XX
  }
}
```

and broader macroeconomic trends, rather than limiting it to a single perspective. Furthermore, we specified a JSON output format to ensure consistency and ease of downstream processing, as unstructured or random outputs can be difficult to interpret and automate.

The LLM processes the prompt and outputs a categorical stock movement prediction, classifying the next day's price movement as Up, Down, or Same. The categorical output is further transformed into a one-hot vector. For example, "Up" is converted to [1,0,0], "Down" is converted to [0,1,0], and "Same" is converted to [0,0,1]. Furthermore, the LLM also outputs the probability scores for each predicted stock movement, reflecting its confidence in the prediction. For example, if the probability is [0.6,0.2,0.2], it means that the "Up" movement has a probability of 0.6, "Down" movement has a probability of 0.2, and "Same" also has a probability of 0.2. These probabilities are essential as they offer a nuanced representation of market uncertainty rather than a binary directional forecast. By transforming textual and numerical data into structured movement probabilities, the LLM effectively distills financial news sentiment into quantifiable features that can be used in stock price prediction models.

The LLM used in this study is "deepseek-r1-distill-llama-70b," which is developed by the American AI company, Groq AI². Users can use this version of the model with the API, but fine-tuning is not available. However, this model has been proven to be able to generate competitive results compared to other LLMs such as LLaMA and Gemma in Reference [25] even without fine-tuning. Therefore, we apply this version of DeepSeek to generate the stock movement and its probability. Furthermore, based on the previous study, applying LLMs directly to predict stock prices cannot outperform models with numerical values as input. Therefore, in this study, LLMs are not applied to directly output the final stock prediction.

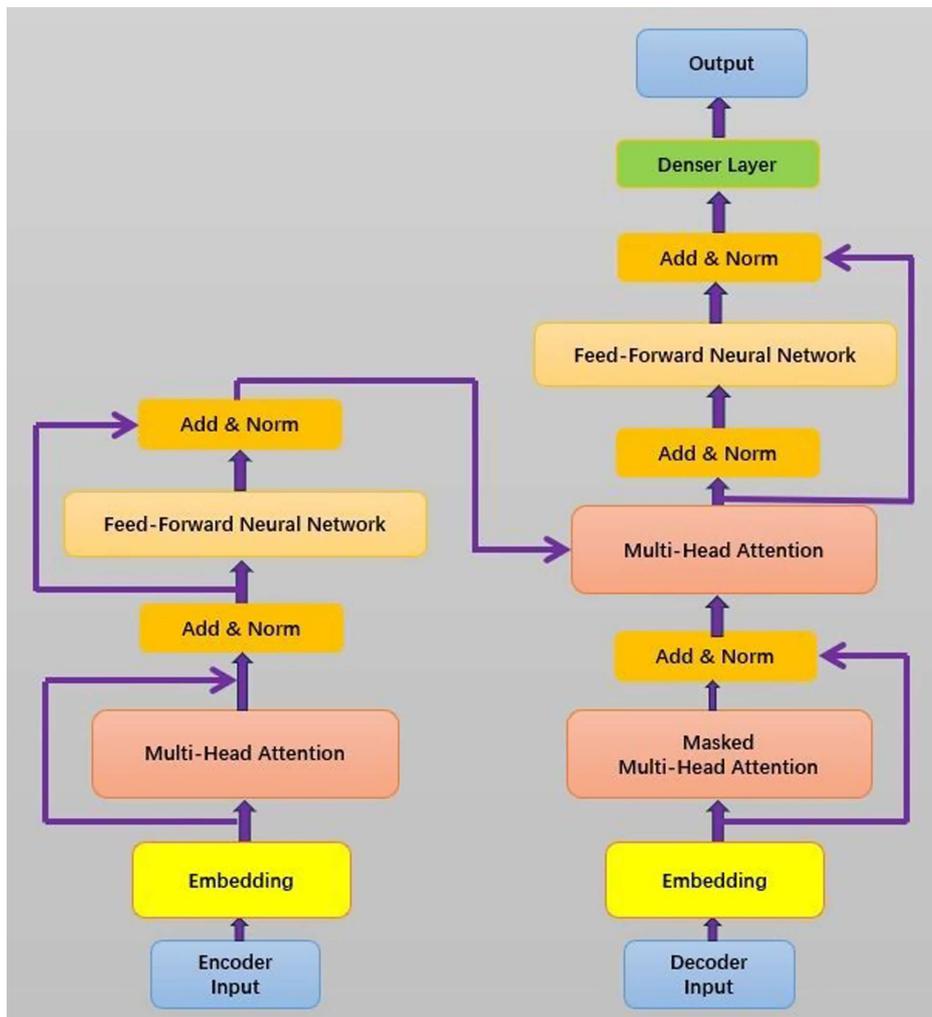
3.3. Transformer-based stock price prediction

The LLM outputs, including stock movement predictions and probability distributions, are integrated with historical stock prices also with a sliding window size of five days to form the input for the

¹<https://eodhd.com/financial-apis/stock-market-financial-news-api>

²<https://groq.com/>

Figure 3
Transformer structure used in this study



Transformer-based stock price prediction model. This model takes advantage of the sequential nature of stock prices and the additional predictive power of sentiment-driven features. Specifically, the input vector consists of the last five days of normalized close prices along with the LLM-generated stock movement. In addition, the probabilities of the Up, Down, and Same predictions are included, ensuring that the model captures the confidence of the LLM in its forecast.

The Transformer model follows an encoder-decoder architecture (Figure 3). The encoder extracts meaningful representations from input features, and the decoder outputs the final stock price prediction. Specifically, the embedding layer first converts the input features into a high-dimensional representation, and the positional encoding is applied to preserve the temporal structure of stock prices. Then, the encoder processes these features through self-attention mechanisms, capturing long-term dependencies and interactions between historical stock prices and LLM-generated predictions. The encoded representation is then passed to the decoder, which utilizes additional self-attention and cross-attention layers to refine the information before generating the final output. Fully connected feed-forward layers further process the decoded representation, enabling the model to predict the next day's stock price by leveraging both historical price movements and sentiment-driven insights.

3.4. Training and evaluation

For training the Transformer model, a regression framework is used, and the objective is to minimize the prediction error between the predicted and actual stock prices. MSE is used as the loss function in the training process. The Adam optimizer is employed with an adaptive learning rate scheduler to improve convergence during training.

For evaluation, we also apply MSE to assess the models. The proposed LLM-Transformer method is compared with various baseline models, including traditional deep learning model such as LSTM, machine learning models such as SVR and RF, and vanilla Transformers without LLM-generated features. By comparing these models, we evaluate the effectiveness of integrating LLM-generated movement predictions and probabilities into the forecasting pipeline. The results provide insights into how financial news sentiment, when processed through an LLM, enhances predictive performance beyond purely historical price-based models.

3.5. Baseline models

All the baseline models only take the historical close price with a sliding window size of five days as input to forecast the close price

on the next day. The following is a brief introduction and model setup of each baseline model.

LSTM: The Long Short-Term Memory (LSTM) is one type of RNN, which is designed to capture sequential dependencies in time-series data and handle the vanishing or exploding gradient problem. It consists of multiple hidden layers and leverages gating mechanisms to retain long-term patterns. In this study, the LSTM is configured with a hidden dimension of 128, two layers, and a sigmoid activation function to handle the output.

TCN: The Temporal Convolutional Network (TCN) is one type of Convolutional Neural Network (CNN) designed to handle sequential data. It leverages dilated causal convolutions to capture long-range dependencies. In this study, the TCN model uses one layer with a kernel size of 2 and a dropout rate of 0.2. The output is passed through a linear layer and a sigmoid activation function for the final prediction.

CNN: The CNN extracts local patterns from time-series data through three convolutional layers with 16, 32, and 64 filters, respectively. Each layer uses ReLU activation, followed by a fully connected layer for regression.

CNN-LSTM: The CNN-LSTM is a hybrid model that combines the strengths of both CNN and LSTM. In this study, the CNN module applies three layers of convolution, while the LSTM handles sequential patterns with a hidden size of 128. The final output is generated using a fully connected layer.

Support Vector Regressor (SVR): In this study, the SVR model employs a linear kernel. The regularization parameter C is set to 10, balancing margin width and misclassification penalty.

Random Forest (RF): In this study, the RF model aggregates the predictions of 300 decision trees to reduce variance and improve accuracy. With a maximum depth of 30, a minimum of two samples per leaf, and a minimum of ten samples for splitting, the model effectively handles complex data patterns while avoiding overfitting.

Vanilla Transformer: The Transformer model adopts an encoder-decoder architecture to predict stock prices. The encoder is composed of three layers with a hidden dimension of 512 and eight attention heads. Positional encoding is integrated to preserve the temporal order of the input. The decoder consists of two layers with the same configuration and generates the next-day stock price by attending to both the encoded representations and the previous outputs.

XGBoost: XGBoost is a gradient boosting algorithm. It also uses decision trees as its base learners, combining them sequentially to improve the model’s performance. In addition, each tree is trained to correct the errors made by the previous tree.

4. Results

The results of the experiments are shown in Table 1. The results provide a comprehensive evaluation of different models for stock price prediction with MSE as the evaluation metric. The models evaluated range from traditional machine learning algorithms, such as RF and SVR, to deep learning approaches such as LSTM, TCN, CNN, and CNN-LSTM, along with Transformer-based architectures.

The findings highlight the superior performance of the proposed Transformer model enhanced with LLM-generated features. Notably, this model consistently achieves the lowest MSE across all target stocks, which suggests that the proposed model demonstrates a strong ability to capture both market dynamics and sentiment-driven signals. In addition, compared to traditional deep learning models, such as LSTM, TCN, and CNN, as well as the hybrid model CNN-LSTM, the proposed approach exhibits a substantial reduction in error, which suggests that integrating LLM-generated features effectively enhances the model’s ability to identify subtle patterns and trends within financial data.

Although the vanilla Transformer already outperforms traditional machine learning models, the inclusion of LLM features further amplifies its predictive power. This improvement can be attributed to the LLM’s ability to extract sentiment information from textual data, which complements the Transformer’s strength in modeling sequential patterns.

Furthermore, the hybrid model, CNN-LSTM, shows competitive performance compared to CNN but fails to outperform LSTM most of the time. However, all of these models still fall short when compared to transformer-based approaches, indicating that leveraging attention mechanisms and sequential encoding plays a critical role in enhancing prediction accuracy. In addition, traditional machine learning models such as RF and SVR exhibit better performance than deep learning models such as LSTM, CNN, and TCN, which suggests that sometimes, simple models are good enough to generate good performance.

In general, the consistent good performance of the proposed Transformer with LLM features highlights the importance of integrating alternative data sources, such as financial news and market sentiment, into stock price prediction models, which is in line with earlier research suggesting that sentiment analysis contributes positively to stock price prediction [4]. However, the results in this study go even further by showing that LLM-generated probabilistic sentiment features offer a richer and more interpretable input to predictive models. Furthermore, by leveraging the complementary strengths of LLMs and Transformers, this approach not only reduces prediction errors but also provides a more holistic

Table 1
Comparison of model performance in terms of MSE for different stocks (MSE)

Model	Apple	HSBC	Tencent	Toyota
LSTM	0.0061	0.0081	0.0077	0.0031
TCN	0.0059	0.0044	0.0039	0.0016
CNN	0.0144	0.0087	0.0094	0.0056
CNN-LSTM	0.0124	0.0099	0.0074	0.0055
Random Forest	0.0055	0.0009	0.0010	0.0004
SVR	0.0063	0.0022	0.0019	0.0015
XGBoost	0.0064	0.0010	0.0043	0.0047
Vanilla Transformer	0.00021	0.00027	0.00032	0.00025
Proposed Transformer with LLM Features	0.00015	0.00020	0.00022	0.00019

understanding of market behavior. These results pave the way for future research in developing more sophisticated models that combine advanced LLMs with deep learning architectures for enhanced financial forecasting.

5. Ablation Study

To further evaluate the individual contribution of the LLM-generated features in stock price prediction, we conduct an ablation study by evaluating different variants of the proposed Transformer model. Specifically, we examine four configurations: (1) a Transformer model without LLM-generated features, using only historical stock prices as input, (2) a Transformer model incorporating only LLM-generated stock movement predictions (up, Down, or Same) without probability estimates, (3) a Transformer model incorporating only LLM-generated movement probabilities without predicted stock movement labels, and (4) the full Transformer model integrating both LLM-generated movement predictions and their corresponding probabilities.

Table 2
Ablation study on the impact of LLM-generated features (MSE), Apple stock price

Model Variant	MSE
Transformer (No LLM Features)	0.00021
Transformer + LLM Movement Prediction Only	0.00019
Transformer + LLM Probability Only	0.00019
Transformer + LLM Movement Prediction + Probability (Full Model)	0.00015

Table 2 shows the performance with different configurations. Generally speaking, the performances between all these four configurations don't have significant differences. However, incorporating LLM-generated features can improve performance. The baseline Transformer model, which relies solely on historical stock prices, exhibits the highest error, indicating that predictions that use historical prices only are slightly less effective in capturing stock market movement. The model that integrates stock movement predictions without probabilities and a model that incorporates probability only achieve the same performance, suggesting that categorical sentiment-based predictions or probability alone provide useful but limited information. However, the most significant performance improvement is observed in the full Transformer model, where both movement predictions and probability estimates are utilized, confirming that the complementary nature of categorical sentiment labels and probabilistic confidence enhances the predictive capability of the model.

These findings highlight the effectiveness of integrating LLM-generated sentiment-driven features into stock price forecasting. The results suggest that while movement predictions provide directional guidance, probability estimates help quantify uncertainty, leading to a more robust and accurate forecasting framework. This ablation study underscores the value of incorporating probabilistic market sentiment alongside traditional stock price features, demonstrating that the LLM-generated features have the potential to improve the model's predictive performance.

6. Conclusion

In conclusion, the integration of LLM-guided market movement signals with Transformer networks demonstrates superior

performance in stock price prediction. The proposed Transformer with LLM Features consistently achieves the lowest MSE across multiple stocks, outperforming traditional machine learning models and even the vanilla Transformer. This success highlights the potential of leveraging LLMs to enhance feature extraction and improve predictive accuracy. Future work can explore further enhancements, such as fine-tuning LLMs for financial contexts and incorporating additional market indicators to enhance robustness and generalization. In addition, different sliding window sizes for the input features will be examined to explore their impact on model performance and to further optimize the prediction accuracy. Moreover, efforts will be made to compare the findings with previous studies by aligning experimental conditions, such as data sources, time periods, and evaluation metrics.

Ethical Statement

This study does not contain any studies with human or animal subjects performed by the author.

Conflicts of Interest

The author declares that he has no conflicts of interest to this work.

Data Availability Statement

The stock data that support the findings of this study can be downloaded from Yahoo Finance using the Python library `finance`. The news data can be downloaded using the Financial News Feed and Stock News Sentiment data API (<https://eodhd.com/financial-apis/stock-market-financial-news-api>).

Author Contribution Statement

Qizhao Chen: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, Writing – review & editing, Visualization.

References

- [1] Ricky, A., Genesius, K., Setiawan, M., Jaya, R. N., & Suhartono, D. (2022). Stock market prediction using support vector regressor with radial basis function kernel. In *Proceedings of the 2022 5th International Conference on Image and Graphics Processing*, 380–387. <https://doi.org/10.1145/3512388.3512444>
- [2] Chhajer, P., Shah, M., & Kshirsagar, A. (2022). The applications of artificial neural networks, support vector machines, and long-short term memory for stock market prediction. *Decision Analytics Journal*, 2, 100015. <https://doi.org/10.1016/j.dajour.2021.100015>
- [3] Cao, R., Liang, X., & Ni, Z. (2012). Stock price forecasting with support vector machines based on web financial information sentiment analysis. In *Advanced Data Mining and Applications: 8th International Conference*, 527–538. https://doi.org/10.1007/978-3-642-35527-1_44
- [4] Chen, Q., & Kawashima, H. (2024). Stock price prediction using LLM-based sentiment analysis. In *2024 IEEE International Conference on Big Data*, 4846–4853. <https://doi.org/10.1109/BigData62323.2024.10825946>

- [5] Du, S., Hao, D., & Li, X. (2022). Research on stock forecasting based on random forest. In *2022 IEEE 2nd International Conference on Data Science and Computer Application*, 301–305. <https://doi.org/10.1109/ICDSCA56264.2022.9987903>
- [6] Vijn, M., Chandola, D., Tikkiwal, V. A., & Kumar, A. (2020). Stock closing price prediction using machine learning techniques. *Procedia Computer Science*, 167, 599–606. <https://doi.org/10.1016/j.procs.2020.03.326>
- [7] Gumelar, A. B., Setyorini, H., Adi, D. P., Nilowardono, S., Latipah, Widodo, A., ..., & Christine, F. (2020). Boosting the accuracy of stock market prediction using XGBoost and long short-term memory. In *2020 International Seminar on Application for Technology of Information and Communication*, 609–613. <https://doi.org/10.1109/iSemantic50169.2020.9234256>
- [8] Hu, Y., Shao, L., La, L., & Hua, H. (2021). Using investor and news sentiment in tourism stock price prediction based on XGBoost model. In *2021 IEEE/ACIS 6th International Conference on Big Data, Cloud Computing, and Data Science*, 20–24. <https://doi.org/10.1109/BCD51206.2021.9581619>
- [9] Deng, S., Huang, X., Zhu, Y., Su, Z., Fu, Z., & Shimada, T. (2023). Stock index direction forecasting using an explainable eXtreme gradient boosting and investor sentiments. *The North American Journal of Economics and Finance*, 64, 101848. <https://doi.org/10.1016/j.najef.2022.101848>
- [10] Sang, S., & Li, L. (2024). A novel variant of LSTM stock prediction method incorporating attention mechanism. *Mathematics*, 12(7), 945. <https://doi.org/10.3390/math12070945>
- [11] Quadir, A., Kapoor, S., Chris Junni, A. V., Sivaraman, A. K., Tee, K. F., Sabireen, H., & Janakiraman, N. (2023). Novel optimization approach for stock price forecasting using multi-layered sequential LSTM. *Applied Soft Computing*, 134, 109830. <https://doi.org/10.1016/j.asoc.2022.109830>
- [12] Nadif, M. A., Rahman Samin, T., & Islam, T. (2022). Stock market prediction using long short-term memory (LSTM). In *2022 Second International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies*, 1–6. <https://doi.org/10.1109/ICAECT54875.2022.9807655>
- [13] Ji, Y., Luo, Y., Lu, A., Xia, D., Yang, L., & Liew, A. W.-C. (2024). Galformer: A transformer with generative decoding and a hybrid loss function for multi-step stock market index prediction. *Scientific Reports*, 14(1), 23762. <https://doi.org/10.1038/s41598-024-72045-3>
- [14] Zhou, L., Zhang, Y., Yu, J., Wang, G., Liu, Z., Yongchareon, S., & Wang, N. (2025). LLM-augmented linear transformer–CNN for enhanced stock price prediction. *Mathematics*, 13(3), 487. <https://doi.org/10.3390/math13030487>
- [15] Muhammad, T., Aftab, A. B., Ibrahim, M., Ahsan, M. M., Muhu, M. M., Khan, S. I., & Alam, M. S. (2023). Transformer-based deep learning model for stock price prediction: A case study on Bangladesh stock market. *International Journal of Computational Intelligence and Applications*, 22(3), 2350013. <https://doi.org/10.1142/S146902682350013X>
- [16] Duan, C., & Ke, W. (2024). Advanced stock price prediction using LSTM and informer models. *Journal of Artificial Intelligence General Science*, 5(1), 141–166. <https://doi.org/10.60087/jaigs.v5i1.183>
- [17] Ren, S., Wang, X., Zhou, X., & Zhou, Y. (2023). A novel hybrid model for stock price forecasting integrating encoder forest and informer. *Expert Systems with Applications*, 234, 121080. <https://doi.org/10.1016/j.eswa.2023.121080>
- [18] Kirtac, K., & Germano, G. (2024). Sentiment trading with large language models. *Finance Research Letters*, 62, 105227. <https://doi.org/10.1016/j.frl.2024.105227>
- [19] Batra, H., Punn, N. S., Sonbhadra, S. K., & Agarwal, S. (2021). BERT-based sentiment analysis: A software engineering perspective. In *Database and Expert Systems Applications: 32nd International Conference*, 138–148. https://doi.org/10.1007/978-3-030-86472-9_13
- [20] Prottasha, N. J., Sami, A. A., Kowsher, M., Murad, S. A., Bairagi, A. K., Masud, M., & Baz, M. (2022). Transfer learning for sentiment analysis using BERT based supervised fine-tuning. *Sensors*, 22(11), 4157. <https://doi.org/10.3390/s22114157>
- [21] Kang, J.-W., & Choi, S.-Y. (2025). Comparative investigation of GPT and FinBERT's sentiment analysis performance in news across different sectors. *Electronics*, 14(6), 1090. <https://doi.org/10.3390/electronics14061090>
- [22] Iacovides, G., Konstantinidis, T., Xu, M., & Mandic, D. (2024). FinLlama: LLM-based financial sentiment analysis for algorithmic trading. In *Proceedings of the 5th ACM International Conference on AI in Finance*, 134–141. <https://doi.org/10.1145/3677052.3698696>
- [23] Xie, Q., Han, W., Zhang, X., Lai, Y., Peng, M., Lopez-Lira, A., & Huang, J. (2023). PIXIU: A large language model, instruction data and evaluation benchmark for finance. In *Proceedings of the 37th Conference on Neural Information Processing Systems*, 33469–33484.
- [24] Wang, N., Yang, H., & Wang, C. D. (2023). FinGPT: Instruction tuning benchmark for open-source large language models in financial datasets. In *Workshop on Instruction Tuning and Instruction Following at NeurIPS 2023*.
- [25] Chen, Q. (2025). Stock price change prediction using prompt-based LLMs with RL-enhanced post-hoc adjustments. In *Proceedings of the 2025 4th International Conference on Bigdata Blockchain and Economy Management*, 195, 475–483. https://doi.org/10.2991/978-94-6463-742-7_46

How to Cite: Chen, Q. (2025). Stock Price Prediction with LLM-Guided Market Movement Signals and Transformer Model. *FinTech and Sustainable Innovation*, 1, A12. <https://doi.org/10.47852/bonviewFSI52025703>