

FB-GAN: A Novel Neural Sentiment-Enhanced Model for Stock Price Prediction

Jainendra Kumar Jain, Ruchit Agrawal

Department of Computer Science, Department of Computer Science
University of Birmingham, University of Birmingham
Birmingham, United Kingdom, Dubai, United Arab Emirates
jkj293@alumni.bham.ac.uk, r.r.agrawal@bham.ac.uk

Abstract

Predicting stock prices remains a significant challenge in financial markets. This study explores existing stock price prediction systems, identifies their strengths and weaknesses, and proposes a novel method for stock price prediction that leverages a state-of-the-art neural network framework, combining the BERT language model for sentiment analysis on news articles and the GAN model for stock price prediction. We introduce the FB-GAN model, an ensemble model that leverages stock price history and market sentiment score for more accurate stock price prediction and propose effective strategies to capture the market sentiment. We conduct experiments on stock price prediction for five major equities (Amazon, Apple, Microsoft, Nvidia, and Adobe), and compare the performance obtained by our proposed model against the existing state-of-the-art baseline model. The results demonstrate that our proposed model outperforms existing models across the five major equities. We demonstrate that the strategic incorporation of market sentiment using both headlines as well as summaries of news articles significantly enhances the accuracy and robustness of stock price prediction.

Keywords: Stock Price Prediction, Sentiment Analysis, GAN, NLP for Finance, BERT, Opinion Mining

1. Introduction

Accurate stock price prediction is a crucial challenge amidst rapid information transmission and complex market dynamics. Traditional quantitative models, while somewhat effective, often fail to grasp market sentiment nuances, relying heavily on historical data. The rise of social media, financial news sites, and online forums has revolutionised the accessibility to stock market-related information. Consequently, investor sentiment, characterized by emotions, opinions and beliefs has emerged as a dynamic force capable of swiftly altering market trends. Based on recent studies done by Xiadong Liu and Li (2023) and Marshan et al. (2023), the qualitative aspects of investor sentiment profoundly impact market movements, affecting the desired rate of return of the investors. By harnessing the power of Natural Language Processing (NLP), Deep Learning models can parse and comprehend vast amounts of textual data generated daily, and gauge the collective sentiment of market participants (Sidogi et al., 2021).¹

Current Machine Learning approaches for stock price prediction primarily rely on autoregressive models such as LSTMs or RNNs (Selvin et al., 2017), (Heaton et al., 2017). The application of modern deep learning approaches for stock price prediction has been limited to the use of the Generative Adversarial Network

(GAN) proposed by Goodfellow (2016) or the Bidirectional Encoded Representations from Transformers (BERT) model proposed by Devlin et al. (2019), without extensively exploring the market sentiment for the current stock (Lin et al., 2021).

In this research, we propose a novel neural method for stock price prediction called FB-GAN, which not only relies on historical stock price data, but also leverages the market sentiment for the particular stock in a strategic manner. We explore and assess the performance of our sentiment-enhanced stock price prediction model using multiple strategies for capturing the market sentiment.

The predictive accuracy and performance of the proposed model FB-GAN is rigorously evaluated and compared against contemporary stock price prediction models using appropriate metrics such as RMSE. We demonstrate that our model strategically incorporates market sentiment data along with historical stock prices and outperforms contemporary approaches for stock price prediction.²

The major contributions of this paper are summarized below:

- We propose a robust neural framework called FB-GAN based on the BERT and GAN models, which leverages market sentiment in a strategic manner along with stock price history for the prediction of upcoming stock prices.
- We conduct experiments using three different

¹Manuscript accepted for publication at FinNLP, LREC-COLING 2024 (<https://lrec-coling-2024.org/>).

²The code will be shared publicly upon acceptance.

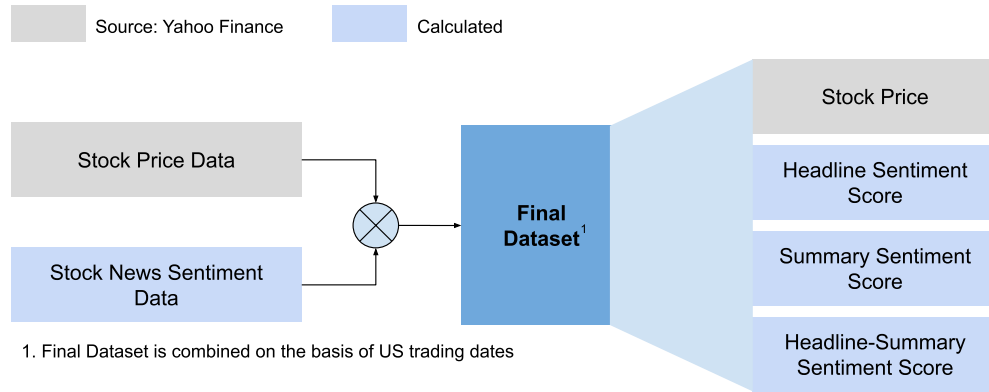


Figure 1: Merging stock price and sentiment score data

strategies to integrate sentiment information from news articles with our stock price prediction model, namely (i) headline, (ii) summary, and (iii) headline and summary combined.

- We demonstrate that the proposed model FB-GAN outperforms contemporary approaches and the combination of headline and summary of the news articles yields the best results for stock price prediction.

2. Related Work

In this section, we briefly review the state-of-the-art techniques for stock price prediction and highlight their limitations, setting the context of our work.

The stock price prediction task dates back to the 1960s, wherein traditional time series analysis methods were used to capture the serial correlation in stock prices (Fama, 1965). These methods, however, often assume stationarity and are not capable of capturing the complexities of time series data.

The advancements in Machine Learning led to the exploration of neural models such as LSTMs (Hochreiter and Schmidhuber, 1997) and CNNs (LeCun et al., 2015) for stock price prediction. (Mehtab and Sen, 2020), (Heaton et al., 2017), (Selvin et al., 2017). Chung et al. (2014) demonstrated that Gated Recurrent Units (GRU) supplemented LSTM networks, which accelerated the training and mitigated the problem of overfitting. Heaton et al. (2017) suggested that the LSTM neural network could be used as an oscillator and were among the first approaches to demonstrate that deep neural networks can detect patterns in financial data. However, LSTMs may have difficulty distinguishing between meaningful patterns and random noise, especially when the data exhibits high volatility or irregular patterns.

While autoregressive models such as LSTMs and RNNs have been explored extensively for stock price prediction, the application of modern neural network architectures for this task remains relatively unexplored. Lin et al. (2021) explored the usage of Generative Adversarial Networks (GANs) and proposed WGAN-GP, an improved GAN model, to make accurate stock price predictions. However, while the WGAN-GP model yields better performance than the previous model, GAN, it only leverages historical stock prices and is unable to capture additional information such as market sentiments.

Akita et al. (2016) explored the usage of LSTMs to incorporate sentiment analysis for stock price prediction. While this method demonstrated the importance of sentiment analysis for this task, it is based on an outdated neural architecture and is unable to capture the market sentiment in a strategic manner. Devlin et al. (2019) introduced BERT, a Transformer-based language model, greatly impacting a number of NLP tasks. BERT has gained popularity for sentiment analysis, extracting valuable insights from news articles, social media, and financial reports.

With the advent of language models such as BERT, research has been conducted to use sentiment analysis on social media and news data for stock price prediction (Weng et al., 2022; Sidogi et al., 2021). These methods are similar to the one proposed by Akita et al. (2016), and employ LSTM with BERT to predict stock prices based on historical prices, with sentiment analysis done on the news headlines of a set of chosen stocks. While these approaches employ the headlines of the news articles for stock price prediction, they fail to capture the entire sentiment of the news articles and are based on the sub-optimal LSTM framework for the time-series prediction.

We provide substantial arguments that sentiment




Figure 2: Proposed Model: FB-GAN

analysis done only on headlines can be misleading and result in poor stock price predictions. In addition to highlighting that headlines do not convey the entire sentiment of the news article, we also propose a neural model which improves stock price prediction by leveraging both historical data as well as market sentiment information, which is done by capturing the entire sentiment of the news article in a strategic manner.

3. Methodology

This section introduces the data collection, data preprocessing, feature engineering, experimental setup, study of existing stock price prediction models and our proposed model.

3.1. Data Collection and Preprocessing

For this study, we selected five stocks: Amazon, Apple, Microsoft, Nvidia and Adobe for stock price prediction based on their 5-year stock price history and market sentiments. The data collection was done in two phases for this project. In the first phase, we gathered news articles related to a particular stock, and in the second phase, we collected the historical price history of the stocks. The news articles related to a particular stock were collected

using the Alpha Vantage API.³

We conduct this study only with high-quality news articles from trustworthy sources. We employ the publicly available news aggregator Alpha Vantage which provides high-quality news articles published by renowned publishers such as The Wall Street Journal, Financial Times, Motley Fool, MarketWatch, etc. We extract information from both the headline and a summary of the news articles, which are essential data points to study the performance of stock price prediction. We extract news articles published during the period 01 Mar 2022 to 31 Jul 2023⁴. The statistics of the data used for our experiments are mentioned in Table 1. The dataset is split randomly into the training and testing sets, such that 80% of the samples are employed for training, and the remaining 20% are used for testing the models.

Historical price data of the stocks was collected using Yahoo Finance's python package yfinance, which gave us data related to a particular stock's close price, open price, high, low and volume for the given time frame. The historical price history collected is from 01 Aug 2018 to 31 Jul 2023.

After performing Exploratory Data Analysis (EDA), we sanitized our dataset to ensure we didn't

³Available at <https://www.alphavantage.co/>

⁴Alpha Vantage does not contain articles published before 01 Mar 2022

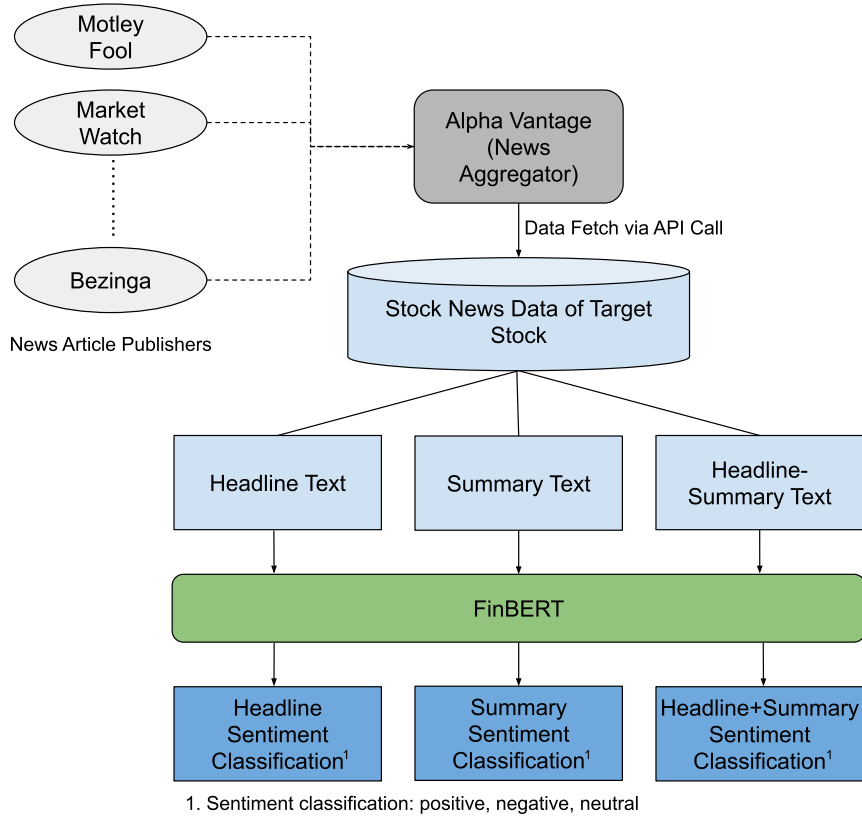


Figure 3: Sentiment Analysis Diagram Flow

Stock Name	Total Articles	Training (80%)	Testing (20%)
Amazon	10.2K	8.2K	2K
Apple	13.8K	11.1K	2.7K
Microsoft	27.6K	22.1K	5.5K
Nvidia	10.8K	8.6K	2.2K
Adobe	1.5K	1.2K	0.3K

Table 1: Count of the news articles captured for each stock from 01 Mar 2022 - 31 Jul 2023 (in thousands)

have any duplicate news articles in our dataset during the data collection process of news articles from the Alpha Vantage API. While performing EDA, it was observed that different news articles vary in terms of their relevance for the stock price prediction task, and some articles could be irrelevant. Hence, in order to ensure that our sentiment analysis is accurate we employ the relevance score provided by Alpha Vantage, which is a measure of how relevant a news article is to a certain stock.

To develop the final dataset used for our experiments, we used the stock price from Yahoo Finance and sentiment scores for each day and combined them based on US Trading dates, as illustrated in Figure 1. The details about the computation of the

sentiment scores and strategies to incorporate the same with the stock price history are presented in Section 3.2. While combining the stock price and sentiment data, we assumed market sentiment for a particular day would have an effect on the next day’s closing price. To handle the dates with no news articles, we have assumed the sentiment for those dates to be neutral i.e., 0 sentiment score. We pass stock news information of all three types: headline, summary, and headline+summary

3.2. Proposed Model: FB-GAN

Our proposed model, FB-GAN is inspired by WGAN-GP which incorporates market sentiment generated by FinBERT for stock price prediction. FB-GAN has two major components, i.e. the generator and the discriminator. The generator is made up of three GRU Units having 1024, 512 and 256 neurons in the three layers respectively; each layer has a dropout ratio of 0.2, followed by three dense layers. The discriminator is made up of three 1-dimensional Convolutional Neural Networks having 32, 64 and 128 neurons in the three layers respectively, with a flattened layer followed by three dense layers and, finally, the output layer, which used linear activation function. The architecture of the proposed model, FB-GAN, is presented in Figure 2.

As shown in Figure 2, the generator transforms random noise along with sentiment scores as an input into data samples that are indistinguishable from real stock price data. The generator aims to produce data which is realistic enough to fool the discriminator. We feed the discriminator with two sample data i.e. real data and generated data. The discriminator aims to classify the real and fake data correctly. The generator and discriminator work in an adversarial manner, where each one tries to outperform the other. Our proposed model, FB-GAN, is trained on 7 features: Adj. Close, High, Low, Close, Open, and Market Sentiment Score. The market sentiments are fed to the neural network as a latent input (co-variant) along with other inputs.

To categorise a piece of news into a particular category, we performed sentiment analysis on each news article using a language model specialized on financial data known as FinBERT. FinBERT (Araci, 2019) is a pre-trained BERT model fine-tuned for financial sentiment classification. FinBERT analyses a textual input and provides an output between 0 and 1 and the sentiment label: positive, negative and neutral. A higher score indicates a higher confidence in the label. In order to assess and analyse the performance of the type of market sentiment information on the stock price prediction, we conduct ablation studies using three different strategies to compute the sentiment scores:

- Using the headline of the article
- Using the summary of the article
- Using the headline and summary of the article

Each of the above are passed to the FinBERT model to obtain the category label and the confidence score for the given news article, as shown in Figure 3. We scaled the sentiment score, obtained from FinBERT by the relevance score of each article obtained from Alpha Vantage for a true estimation of the overall sentiment of each news article. Since neural networks can only process numerical input, we pre-process the data before feeding it to the network. We feed two types of inputs to the neural network i.e. stock price data and market sentiment data. The stock price data is already in numerical form; however, the output from FinBERT sentiment classification is in the form of textual labels, namely, positive, negative and neutral. In order to transform it to numerical form, we assign positive articles a value of 100, negative articles a value of -100 and neutral articles a value of 0. To calculate the sentiment score for a particular day, we define the Sentiment Score SS_n as follows:

$$SS_n = \frac{\sum_{i=1}^N (CS_{pos} \times 100) + \sum_{j=1}^M (CS_{neg} \times -100)}{N + M + P} \quad (1)$$

where N , M , and P represent the total number of positive, negative and neutral articles for a particular day, respectively, CS_{pos} (Confidence score - positive) represents the confidence score of the positive article(s), CS_{neg} (Confidence score - negative) represents the confidence score of the negative article(s).

The scores computed using these mechanisms are then fed to our model alongside the stock price history to perform the prediction of the upcoming stock prices. The optimizer used is Adam, with a learning rate of 0.000128, the number of epochs is 160, and a batch size of 128.

4. Experimental Setup

To conduct the experiments, we employ the Python 3 Google Compute Engine. The hardware setup includes a Nvidia Tesla T4 GPU with a 2-core Intel Xenon CPU 2.2 GHz, supported by 13GB RAM and 16 GB GPU Memory. 80% of the samples are used for training, and 20% are used for testing the model. The models are implemented using the deep learning framework Keras, with a Tensorflow backend.

We compare the performance of our model with the following existing approaches:

- Vanila RNN model: The predictions are done based on the Adjusted close price as the input feature (3 days of Adjusted close price to predict the Adjusted close price of the next day). The RNN model comprises 5 layers: 1 input layer, 3 hidden layers, and 1 output layer. The optimizer used is Root Mean Square Propagation (RMSprop), the loss function is the Mean Squared Error (MSE), the number of epochs is set to 100, and the batch size is set to 150.
- LSTM: In the LSTM model, we use a similar input vector as we did in the case of Vanilla RNN, where we use 3 days of Adjusted close price to predict the Adjusted close price of the next day. The LSTM model contains 5 layers: 1 input layer, 3 hidden layers, and 1 output layer. The optimizer is Adam, the loss function is Mean Squared Error (MSE), the number of epochs is set to 100, and the batch size is set to 150. A dropout layer is added after each LSTM layer to prevent overfitting. The dropout ratio is set to 0.2.
- GAN: The generator uses a three-dimensional array of tensors, time steps, and features, similar to the vanilla RNN. The model GAN is trained on 6 features: Adj. Close, Open, High, Low, Close and Volume, using 3-time steps to give the prediction of the next day's Adj. close price. The optimizer used is Adam, with a

	Amazon	Apple	Microsoft	Nvidia	Adobe	Average
Vanilla RNN	5.30	9.34	16.24	25.98	16.66	14.71
LSTM	4.31	<u>6.53</u>	9.44	28.19	15.76	12.85
GAN	4.49	12.73	16.74	23.02	17.76	14.95
WGAN-GP	5.03	6.98	18.29	18.30	14.67	12.65
FB-GAN (Headline)	4.78	7.61	12.53	<u>15.58</u>	21.10	<u>12.32</u>
FB-GAN (Summary)	<u>4.30</u>	8.13	12.26	19.01	21.67	13.07
FB-GAN (Headline+Summary)	4.01	4.35	<u>10.08</u>	14.19	<u>15.73</u>	9.67
Average	4.52	7.49	13.00	20.21	17.21	–

Table 2: Comparison of results of different models based on RMSE. Best performing model highlighted in bold, second best performing model underlined




Figure 4: FB-GAN (Headline+Summary) Actual vs. Predicted Stock Price Graph of Amazon

learning rate of 0.00016, the model is trained for 165 epochs with a batch size of 128. Leaky Rectified Linear Unit (ReLU) is used as an activation function among all layers except the output layer, which is a sigmoid activation function. The model is tuned with a learning rate between 0.0003, number of epochs of 300 and a batch size between 64 to 512.

- WGAN-GP: The architecture of WGAN-GP is based on the GAN model; however, the output layer of the discriminator of the WGAN-GP is a linear activation function instead of a sigmoid function, and an additional gradient penalty is added to the discriminator. The optimizer used is Adam, with a learning rate of 0.000115. The model is trained for 100 epochs, with a batch size of 128. The discriminator and generator are the same as the basic GAN; however, the discriminator is trained once, and the generator is trained thrice.

5. Results and Discussions

5.1. Quantitative Analysis

We compare the performance obtained by our proposed model (FB-GAN) with five existing neural network baseline models. The results obtained by the models are presented in Table 2. We used Root Mean Square Error (RMSE) as the evaluation metrics, defined as Equation 2:

$$\text{RMSE}(y, \hat{y}) = \sqrt{\frac{\sum_{i=0}^{N-1} (y_i - \hat{y}_i)^2}{N}} \quad (2)$$

where y_i is the actual (true) value of the i^{th} data point, \hat{y}_i is the predicted value of the i^{th} data point and N is the total number of data points. A lower RMSE value signifies a better model as the predicted values are as close as possible to the target values.

Data Fields	News Article #1	News Article #2	News Article #3
Ticker	AMZN	AMZN	AMZN
Date	25/04/23	07/03/22	10/03/22
Time	15:13:49	16:08:48	13:20:25
Headline	"What's Going On With Amazon Shares - Amazon.com (NASDAQ:AMZN)"	"Why Amazon, Meta Platforms And Microsoft Shares Are Falling Today"	"Why Amazon Shares Are Rising"
Summary	"Amazon.com, Inc. AMZN shares are trading lower by 1.96% to \$104.13. The stock is trading lower possibly in anticipation of the company's first-quarter earnings report, confirmed for Thursday's after-hours session."	"Shares of several companies in the broader technology sector, including Amazon.com, Inc. (NASDAQ: AMZN), Meta Platforms Inc (NASDAQ: FB) and Microsoft Corporation (NASDAQ: MSFT), are all trading lower as stocks fall amid the continued escalation of the Russia-Ukraine conflict."	"Amazon.com, Inc. (NASDAQ: AMZN) shares are trading higher by 4.7% at \$2,917.75 after the company reported a 20-for-1 stock split and a \$10 billion share buyback. Amazon says, subject to shareholder approval of the stock split, each company shareholder of record at the close..."
Headline+Summary	"What's Going On With Amazon Shares - Amazon.com (NASDAQ:AMZN) Amazon.com, Inc. AMZN shares are trading lower by 1.96% to \$104.13. The stock is trading lower possibly in anticipation of the company's first-quarter earnings report, confirmed for Thursday's after-hours session."	"Why Amazon, Meta Platforms And Microsoft Shares Are Falling TodayShares of several companies in the broader technology sector, including Amazon.com, Inc. (NASDAQ: AMZN), Meta Platforms Inc (NASDAQ: FB) and Microsoft Corporation (NASDAQ: MSFT), are all trading lower as stocks fall amid the continued escalation of the Russia-Ukraine conflict."	"Why Amazon Shares Are RisingAmazon.com, Inc. (NASDAQ: AMZN) shares are trading higher by 4.7% at \$2,917.75 after the company reported a 20-for-1 stock split and a \$10 billion share buyback. Amazon says, subject to shareholder approval of the stock split, each company shareholder of record at the close..."
Source	Benzinga	Benzinga	Benzinga
Relevance Score	0.9267	0.5502	0.9836
Headline FL ^a	Neutral	Negative	Positive
Headline FS ^b	1	0.8112	1
Summary FL ^a	Negative	Negative	Positive
Summary FS ^b	0.9999	0.9706	0.9903
Headline+Summary FL ^a	Negative	Negative	Positive
Headline+Summary FS ^b	0.9963	0.9913	0.9999

^a FL stands for FinBERT Label

^b FS stands for FinBERT Score

Table 3: Sample News Data after Sentiment Analysis

The results obtained by our proposed model FB-GAN are compared with existing approaches and are presented in Table 2. The effect of stock price prediction with and without sentiment information can be observed by comparing the result obtained by FB-GAN (Headline+Summary) with WGAN-GP, the best-performing stock price prediction model which uses only historical price data. FB-GAN outperforms WGAN-GP by 23.6% in terms of the RMSE value (Table 2, rows 4-6). Additionally, FB-GAN also outperforms other baseline models, namely the RNN, LSTM and GAN models (Table 2,

rows 1-3).

In addition to demonstrating the impact of sentiment, we conduct ablation studies using three different strategies to incorporate information from news articles, namely (i) headline, (ii) summary, and (iii) summary and headline combined. FB-GAN yields the best results based on the sentiment obtained from the headline and summary combined for each stock, with an average RMSE of 9.67, followed by the headline sentiment with an average RMSE of 12.32 and lastly using the summary sentiment with an average RMSE of 13.07.

Figure 4 juxtaposes the actual stock price of Amazon against that predicted by the proposed model FB-GAN on the test data, using the Headline+summary strategy. It can be observed that our predicted stock price mimics the actual close price very closely, demonstrating its efficacy for stock price prediction. Although each stock comprises complex time series data, our FB-GAN model performs well in predicting the stock price of each stock.

5.2. Qualitative Analysis

In addition to providing a quantitative analysis of the results, we also present a qualitative analysis of the results obtained by FB-GAN, and demonstrate the importance of sentiment analysis in stock price prediction by comparing the different strategies to capture market sentiment.

Table 3 presents a sample of our stock news dataset after performing sentiment analysis using FinBERT (refer to Figure 3). In our stock news dataset, we have three text parameters: Headline, Summary and Headline+Summary, which we had passed through FinBERT and obtained the FinBERT Label (FL) and FinBERT Score (FS) for each parameter. The FinBERT Label can be any one of three labels: Positive, Negative and Neutral; and the FinBERT Score can be any value between 0 and 1, where a lower score represents low confidence and higher score represents high confidence.

On comparing the Headline FL, Summary FL and Headline+Summary FL of News Article #1 of Table 3, we observe Headline FL is classified as Neutral, Summary FL is classified as Negative and Headline+Summary is classified as Negative. In general, news article headlines could be incomplete and misleading to attract readers' attention and could lead to incorrect classification when sentiment analysis is performed on them. The headline of News Article #1, *"What's Going On With Amazon Shares"* may spark curiosity in reader; and while a human might delve deeper to understand the topic to make an informed opinion, an ML-model might fail to capture the sentiment based on the headline alone. Previous studies have solely relied on the headlines, for incorporating sentiment analysis in stock price prediction (Sidogi et al., 2021; Weng et al., 2022; Li et al., 2023). In this case, the headline suggests neutrality, but the summary paints a negative picture, with *"Amazon shares trading lower by 1.96% to \$104.13"*. This highlights the limitation of relying solely on headlines or summaries for sentiment analysis, as they may only present half the picture. By combining the headline and summary, FinBERT can accurately classify the article as negative with 99.63% confidence, demonstrating the importance of complete information for accurate

stock price prediction. Similarly, for News Article #2 of Table 3, on comparing the FinBERT Score (FS) obtained after passing Headline, Summary and Headline+Summary through FinBERT, we observe that FS of Headline+Summary is 99.13%, followed by FS of Summary which is 97.06%, followed by FS of Headline which is 81.12%, which proves statistically as well that Headline+Summary provides a higher confidence on the estimated label than its counterparts.

Based on News Article #3 of Table 3, we perform pre-hypothesis testing, where we compare the results from the FinBERT classification with the actual stock price movement. It can be observed that based on the Headline and Summary of the News Article - *"Why Amazon Shares are rising. Amazon [...] shares [...] trading higher by 4.7% at \$2917.75 after the company reported 20-for-1 stock split and a \$10 billion share buyback [...]"*, the News Article is classified as positive with a confidence score of 99.99% and following this news, the stock price shows a bullish (upward) trend for several days.

The qualitative analysis thereby corroborates the finding that our sentiment-enhanced model yields improved performance owing to the correlation between market sentiment and stock price movement. It also confirms that the Headline+Summary combined strategy provides a more accurate estimation of the sentiment than individual strategies, leading to better stock price prediction.

6. Conclusion and Future Work

This paper presents a novel sentiment-enhanced neural model called FB-GAN, and demonstrates that it outperforms existing approaches for stock price prediction. The experimentation validates our hypothesis that integrating market sentiment in a strategic manner using state-of-the-art language models improves the performance of stock price prediction. We demonstrate that the Headline & Summary combined strategy yields the best results for stock price prediction (an improvement of 21.5% and 26% respectively in the average RMSE scores when considering Headline alone and summary alone respectively).

Future directions to improve our proposed model could be inspired from the Efficient Market Hypothesis (EMH), wherein more correlated factors, such as gold prices, bank rates, etc., are leveraged while training the model for stock price prediction. Another possible direction for future work involves modifying our proposed model to consider the real-time stock price and market sentiment data to predict the stock prices which can be used for Intra-day trading.

7. Bibliographical References

- Ryo Akita, Akira Yoshihara, Takashi Matsubara, and Kuniaki Uehara. 2016. [Deep learning for stock prediction using numerical and textual information](#). *2016 IEEE/ACIS 15th International Conference on Computer and Information Science, ICIS 2016 - Proceedings*.
- Dogu Araci. 2019. [FinBERT: Financial Sentiment Analysis with Pre-trained Language Models](#). *CoRR*, abs/1908.10063.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. [Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova Google, and AI Language. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). *Proceedings of the 2019 Conference of the North*, pages 4171–4186.
- Eugene F. Fama. 1965. [The behavior of stock-market prices](#). *The Journal of Business*, 38:34.
- Ian Goodfellow. 2016. *Deep learning / Ian Goodfellow, Yoshua Bengio and Aaron Courville*.
- J. B. Heaton, N. G. Polson, and J. H. Witte. 2017. [Deep Learning for Finance: Deep Portfolios](#). *Applied Stochastic Models in Business and Industry*, 33:3–12.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-Term Memory](#). *Neural Computation*, 9:1735–1780.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature*, 521(7553):436–444.
- J.-B. Li, S.-Y. Lin, F.-Y. Leu, and Y.-C. Chu. 2023. [Stock Price Trend Prediction Using LSTM and Sentiment Analysis on News Headlines](#). *Advances on Broad-Band Wireless Computing, Communication and Applications: Proceedings of the 17th International Conference on Broad-Band Wireless Computing, Communication and Applications (BWCCA-2022)*. *Lecture Notes in Networks and Systems (570)*, pages 282 - 91, Berlin, Germany.
- Hung Chun Lin, Chen Chen, Gao Feng Huang, and Amir Jafari. 2021. [Stock price prediction using Generative Adversarial Networks](#). *Journal of Computer Science*, 17:188–196.
- A. Marshan, M. Mbedzi, and A. Ioannou. 2023. [Exploring the relationship between news articles and stocks market movements: A sentiment analysis perspective](#). *Intelligent Computing: Proceedings of the 2023 Computing Conference*. *Lecture Notes in Networks and Systems (739) vol.2* pages 1167 - 83, Berlin, Germany.
- Sidra Mehtab and Jaydip Sen. 2020. Stock price prediction using CNN and LSTM-based Deep learning models. In *2020 International Conference on Decision Aid Sciences and Application (DASA)*, pages 447–453. IEEE.
- Sreelekshmy Selvin, R Vinayakumar, EA Gopalakrishnan, Vijay Krishna Menon, and KP Soman. 2017. [Stock price prediction using LSTM, RNN and CNN-sliding window model](#). In *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 1643–1647. IEEE.
- Thendo Sidogi, Rendani Mbuva, and Tshilidzi Marwala. 2021. [Stock price prediction using sentiment analysis](#). In *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 46–51.
- Xiaojian Weng, Xudong Lin, and Shuaibin Zhao. 2022. [Stock price prediction based on LSTM and BERT](#). In *2022 International Conference on Machine Learning and Cybernetics (ICMLC)*, pages 12–17.
- Bing Han Xiaodong Liu and Luanfeng Li. 2023. [Impact of investor sentiment on portfolio](#). *Emerging Markets Finance and Trade*, 59(3):880–894.