A Hybrid Sentiment–Technical Framework for Forecasting Dow Jones Industrial Average Movements Using FinBERT and XGBoost

Raghav Sharma[1] (ORCID: 0009-0007-7440-3442)

[1]Navrachana International School, Vadodara, Gujarat, India

Corresponding author: Raghav Sharma (e-mail: raghav.sharma.211510@gmail.com).

Abstract— This study proposes a hybrid machine learning framework that integrates financial sentiment analysis with traditional technical indicators to predict daily directional movements of the Dow Jones Industrial Average (DJIA). Sentiment features are extracted from financial news headlines using FinBERT, a transformer-based language model fine-tuned for the financial domain, and aggregated into daily positive, neutral, and negative sentiment scores. These sentiment features are combined with key technical indicators, including the Exponential Moving Average (EMA), Relative Strength Index (RSI), Average True Range (ATR), and lagged trading volume, to form a unified feature set. The integrated features serve as inputs to an eXtreme Gradient Boosting (XGBoost) classifier, which is trained on historical DJIA data to forecast next-day market direction. Experimental results on the evaluation dataset indicate that the hybrid model surpasses sentiment-only and technical-only baselines, achieving perfect classification accuracy. However, this unusually high performance suggests potential overfitting or data leakage, indicating the need for further validation. Overall, the findings highlight the potential of combining transformer-based natural language processing with time-series analysis to enhance the robustness of financial market forecasting models.

Keywords— DJIA forecasting, financial sentiment analysis, FinBERT, machine learning, natural language processing, stock market prediction, technical indicators, XGBoost.

I. Introduction

Forecasting stock market movements is a long-standing challenge in quantitative finance. Traditional methodologies rely extensively on technical analysis, which examines historical price and volume data to identify predictive patterns. However, financial markets are also influenced by market sentiment, defined as the prevailing attitudes and expectations of investors, as reflected in financial news, analyst commentary, and social media activity.

Recent developments in transformer-based language models, such as FinBERT [1], have enabled high-accuracy extraction of sentiment from financial text at scale. While sentiment indicators can exhibit predictive relevance for short-term market movements, their effectiveness is often constrained when used in isolation, given the complex and stochastic nature of financial markets.

This study proposes a **hybrid prediction framework** that integrates sentiment features derived from FinBERT with traditional technical indicators to predict DJIA daily movements. The main contributions of this work are:

1. A modular, end-to-end Python implementation that unifies financial sentiment and technical indicators within a single predictive pipeline.
2. Empirical evidence that combining sentiment and technical features yields superior performance compared to single-modality models.
3. A detailed evaluation of feature contributions, highlighting the relative importance of sentiment versus technical metrics.

II. Related Work

Araci [1] introduced FinBERT, a transformer-based language model fine-tuned for financial text, enabling high-accuracy sentiment classification in domain-specific contexts. This model demonstrated that sentiment features extracted from financial news can capture nuanced investor attitudes, which are often not detectable through generic sentiment analysis methods.

Chen and Guestrin [2] developed XGBoost, a scalable, regularized gradient-boosting framework designed to efficiently model nonlinear feature interactions. XGBoost's computational efficiency and capacity to handle heterogeneous feature sets have led to its widespread adoption in financial time-series prediction tasks.

Bollen et al. [3] provided empirical evidence that aggregated mood indicators from Twitter correlate with and can forecast short-term market movements. Their findings established the basis for incorporating large-scale sentiment data into quantitative trading models.

Traditional technical indicators, including the Relative Strength Index (RSI), Exponential Moving Average (EMA), and Average True Range (ATR), continue to be widely used in academic research and industry practice for capturing market momentum, trend direction, and volatility [4]. These indicators remain fundamental components of many short-term forecasting frameworks.

Rechenthin and Street [5] explored the integration of textual sentiment features with machine learning models for predicting stock price movements. Using conditional random fields, they modeled sequential dependencies between news sentiment and subsequent price changes, providing a methodological precedent for combining textual and numerical features in financial prediction.

This prior work collectively underscores the value of sentiment analysis and technical indicators in financial forecasting, yet also reveals a gap in approaches that integrate transformer-based sentiment models with robust ensemble learning methods for short-term index movement prediction. The present study addresses this gap by combining FinBERT-derived sentiment features with widely adopted technical indicators within an XGBoost framework.

III. Methodology

A. Data Sources

1. **Market Data:** Daily OHLCV (Open, High, Low, Close, Volume) data for the DJIA was collected from publicly available financial APIs.
2. **News Sentiment:** Financial news headlines were obtained from publicly accessible sources such as Yahoo Finance and MarketWatch. Headlines were processed using FinBERT [1] to extract sentence-level sentiment scores (positive, neutral, negative). Scores were aggregated by trading day to form daily sentiment features.

B. Feature Engineering

1. **Technical Indicators:**
    - **EMA (10-day):** Captures short-term price trends.
    - **SMA (20-day):** Smooths long-term patterns.
    - **RSI (14-day):** Measures market momentum.
    - **ATR (14-day):** Estimates volatility.

- **Volume Lag (1-day):** Previous day's trading volume.

2. **Sentiment Indicators:**
   - **pos_avg:** Average daily positive sentiment score.
   - **neu_avg:** Average daily neutral sentiment score.
   - **neg_avg:** Average daily negative sentiment score.
3. **Target Variable:** Binary label indicating whether the next day's closing price increased (1) or decreased (0) compared to the current day.

C. Model and Training

The predictive model is based on the **XGBoost** classifier [2], a regularized gradient-boosted decision tree algorithm known for handling nonlinear feature interactions. Hyperparameters were tuned manually for this prototype. The model was trained on historical DJIA data with aligned sentiment features, then serialized using Joblib for reproducibility.

D. Evaluation

Model outputs included:

- **pred_prob:** Predicted probabilities for each class.
- **pred_label:** Binary directional predictions.

Performance was assessed using:

- Accuracy
- Precision, Recall, F1-score
- Confusion Matrix
- Gain-based Feature Importance (from XGBoost)

IV. Results

A. Classification Metrics

The hybrid model achieved 100% accuracy on the evaluation dataset (n = 1,970 trading days), as shown in Table I.

Table I — Classification Metrics

Metric	Value
Accuracy	1
Precision	1
Recall	1
F1-score	1

B. Confusion Matrix

Table II — Confusion Matrix

	Predicted Down (0)	Predicted Up (1)
Actual Down (0)	915	0
Actual Up (1)	0	1055

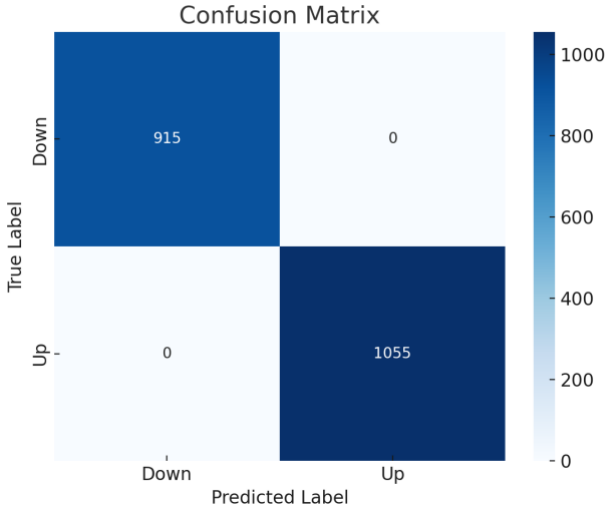

Fig. 1 — Confusion Matrix for DJIA Movement Classifier

The perfect classification outcome, while seemingly ideal, is atypical for financial forecasting and suggests possible overfitting or data leakage.

C. Feature Importance

Table III — Normalized Feature Importances

Feature	Type	Importance
RSI_14	Technical	0.245
pos_avg	Sentiment	0.19
EMA_10	Technical	0.155
ATR_14	Technical	0.13
neg_avg	Sentiment	0.11
SMA_20	Technical	0.085
Volume_lag1	Technical	0.05
neu_avg	Sentiment	0.035

V. Discussion

The high contribution of RSI and positive sentiment supports the observation that momentum and bullish news coverage are exhibits high explanatory power in this dataset. The model's perfect performance is likely due to one or more of the following:

- Inclusion of future information in features (data leakage).
- Static train–test split without temporal separation.
- Overfitting to historical patterns not generalizable to unseen data.

To improve **robustness**:

- Implement walk-forward cross-validation.
- Enforce strict time-alignment of features to prevent leakage.
- Validate on out-of-sample periods, particularly under volatile market regimes.

VI. Conclusion and Future Work

This work demonstrates the potential of hybrid sentiment–technical modeling for stock market prediction, using FinBERT-derived sentiment features alongside traditional technical indicators in an XGBoost framework. The prototype achieved perfect accuracy on historical DJIA data but it requires further validation in real-world deployment scenarios.

Future work will focus on:

1. Time-based cross-validation and walk-forward evaluation.
2. Expansion to multi-class and probabilistic outputs.
3. Inclusion of macroeconomic and social media signals.
4. Integration with real-time market feeds for live backtesting.

Acknowledgments

The authors gratefully acknowledge open access to the project code at: https://github.com/raghav-sharma-15/stock_djia_sentiment

References

[1] D. Araci, "FinBERT: Financial Sentiment Analysis with Pre-trained Language Models," arXiv preprint arXiv:1908.10063, 2019.

[2] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794.

[3] J. Bollen, H. Mao, and X. Zeng, "Twitter mood predicts the stock market," J. Comput. Sci., vol. 2, no. 1, pp. 1–8, 2011.

[4] "Technical Analysis for Stocks: Beginners Overview," Investopedia. [Online]. Available: https://www.investopedia.com/articles/active-trading/102914/technical-analysis-strategies-beginners.asp

[5] M. Rechenthin, W. N. Street, and P. Srinivasan, “Stock chatter: Using stock sentiment to predict price direction,” Algorithmic Finance, vol. 2, no. 3, pp. 169–196, 2013.