TIME-GAN MODEL FOR STOCK PREDICTION

***Sandip D Gurule*[1]*, Dr. G. P.Potdar*[2]**
[1] Student, Dept. of Data Science, Pune Institute of Computer Technology, Pune
[2]Professor, Dept. of Data Science, Pune Institute of Computer Technology, Pune

Abstract—
This model proposes a deep learning framework for stock price prediction using Time-series Generative Adversarial Networks (Time-GAN). Time-GAN uniquely combines autoregressive modeling with adversarial training to generate realistic time-series data while preserving temporal dynamics. We apply this architecture to predict Apple Inc. stock prices, integrating additional macroeconomic indicators such as the S&P 500, NASDAQ Composite, and U.S. Dollar Index. To further enhance predictive accuracy, financial sentiment features are extracted using FinBERT, a transformer model fine-tuned on financial news. The enriched dataset allows Time-GAN to learn both quantitative patterns and market sentiment trends. Performance is evaluated against conventional models, including LSTM and GRU, demonstrating that Time-GAN achieves superior forecasting accuracy and better generalization on unseen data. These results highlight the potential of generative models in financial time-series analysis and their applicability in real-world investment strategies..

Keywords: **Time-GAN, FinBERT, NLP, LSTM, GRU, CNN**

I. INTRODUCTION

Predicting stock prices is a longstanding and challenging problem in financial analysis, primarily due to the volatile, non-linear, and noisy nature of market data. Despite many efforts to model stock price movements, traditional statistical and machine learning methods such as AutoRegressive Integrated Moving Average (ARIMA) and Long Short-Term Memory (LSTM) networks have had limited success in fully capturing the dynamic intricacies of financial markets.

In recent years, Generative Adversarial Networks (GANs) have gained significant attention for their ability to generate realistic synthetic data, particularly in fields like image generation. However, their application to time-series forecasting, especially in financial contexts, is still relatively new. Time-series Generative Adversarial Networks (Time-GAN) offer a solution by combining the advantages of GANs with recurrent architectures, thus preserving both the temporal structure and underlying data distributions. This makes Time-GAN especially well-suited for forecasting tasks that require understanding both sequential dependencies and statistical properties.

This study leverages Time-GAN for stock price prediction, specifically focusing on Apple Inc.'s daily closing prices. Our model is further enhanced by including macroeconomic indicators such as the S&P 500, NASDAQ Composite, and the U.S. Dollar Index. Additionally, we incorporate key technical indicators—such as Moving Average Convergence Divergence (MACD) and Relative Strength Index (RSI)—which help the model better detect price trends and market momentum.

The dataset spans from July 1, 2010, to June 30, 2020, with a seven-year training period and two years reserved for validation. By combining time-series data with financial sentiment, this approach aims to offer a more comprehensive solution for predicting stock price movements compared to traditional methods.

II. PROBLEM STATEMENT

Accurately forecasting stock prices remains a formidable AI challenge, driven by the market's inherent volatility, non-linearity, and the interplay of multiple quantitative and qualitative factors. Conventional time-series methods often struggle to capture long-range temporal dependencies and to integrate external signals such as news sentiment. In this study, we first establish baselines using two recurrent architectures—Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU)—to model sequential price patterns. To address their limitations, we propose a Time-series Generative Adversarial Network (TimeGAN) framework, featuring a GRU-based generator that synthesizes plausible future price trajectories and a CNN-based discriminator that assesses their realism. Furthermore, we incorporate FinBERT-derived sentiment scores from daily financial news to enrich our input feature space, thereby enabling the model to respond to market psychology during major events. This combined approach aims to improve both the accuracy and robustness of stock price predictions in complex, data-rich environments.

III. RELATED WORKS

Recent progress in machine learning and deep learning has significantly influenced the field of stock price forecasting. A wide spectrum of models—ranging from conventional statistical techniques to sophisticated hybrid architectures—has been developed to enhance prediction performance and generalization.

Ilyas et al. [1] introduced a hybrid forecasting approach that utilizes a fully modified Hodrick–Prescott (HP) filter in combination with methods such as Support Vector Regression

(SVR), ARIMA, LSTM, and GRU. Although this approach effectively smooths noisy time-series data, its accuracy deteriorates during periods of high market volatility and lacks adaptability across diverse financial markets. Zhang et al. [2] proposed a GAN-based model tailored for stock market prediction, achieving superior results compared to classical models in terms of RMSE and MAE. Nevertheless, the approach faced challenges such as unstable GAN training and a relatively limited testing scope.

In a separate study, Zhou et al. [3] applied GANs to high-frequency trading scenarios, showcasing their ability to replicate fine-grained market behaviors. However, the model's dependence solely on historical price and volume data limited its capability to incorporate broader economic trends. Bi et al. [4] developed a comprehensive prediction pipeline that included feature engineering and Bayesian optimization. While this framework delivered strong forecasting performance, it encountered scalability bottlenecks and over fitting risks due to its high-dimensional input space.

Beyond finance, machine learning has also proven effective in other domains. Rao et al. [5] applied ML techniques like SVM and MLP to diagnose Alzheimer's Disease using 3D MRI data. Although the models were effective, the limited availability of annotated medical images posed a significant challenge. In another study [6], the same group extended their research into financial time-series by employing Fast RNN and a CNN–BiLSTM hybrid model. The approach showed improvements in accuracy and computational efficiency, though issues with model interpretability and generalization persisted.

Yoon and Jarrett [7] introduced TimeGAN, an innovative framework that merges adversarial learning with supervised time-series modeling. The model effectively generates realistic sequential data by preserving both temporal dependencies and feature distributions. Mozaffari and Zhang [8] explored a different direction, combining ensemble learning with models like Transformers, ARIMA, and Linear Regression. Their stacking ensemble achieved high predictive power but added considerable complexity and reduced model interpretability.

Overall, the literature reflects a trend toward combining multiple data modalities and model types—ranging from GANs and hybrid networks to sentiment-aware systems—each offering distinct advantages and encountering unique limitations. These evolving methodologies highlight the ongoing need for innovative, interpretable, and adaptable approaches in stock price forecasting.

IV. Architecture

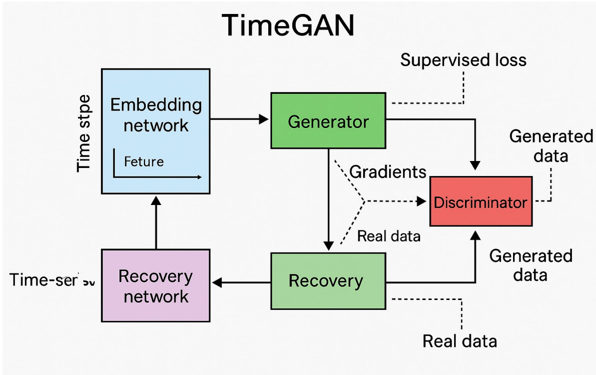

The TimeGAN (Time-series Generative Adversarial Network) architecture is designed to generate realistic and temporally consistent synthetic time-series data by combining the strengths of both recurrent neural networks (RNNs) and generative adversarial networks (GANs).

1) Components of Time GAN:

1. Embedding Network
   - This network maps original time-series data into a latent space.
   - It captures temporal patterns using RNN layers (typically LSTM or GRU).
   - Its purpose is to transform high-dimensional sequential data into a compressed representation.
2. Recovery Network
   - Paired with the embedding network, it reconstructs the original time-series data from the latent space.
   - This forms an auto encoder structure (embedding + recovery), ensuring the latent representation preserves meaningful sequential information.
3. Generator
   - Takes random noise and generates synthetic latent time-series sequences.
   - It also uses RNN layers to capture the temporal dependencies during generation.
4. Discriminator
   - Attempts to distinguish between real and synthetic latent sequences.
   - It learns to recognize whether a sequence comes from real data or from the generator.
5. Supervised Network
   - Predicts the next time step in the latent space.
   - Ensures that the generated sequences maintain logical time progression and temporal relationships.

2)Architectural Flow:

- Original time-series data → Embedding Network → Latent space → Recovery Network → Reconstructed data
- Random noise → Generator → Synthetic latent space → Discriminator
- Supervised Network predicts future latent steps to reinforce temporal accuracy

By combining these modules, TimeGAN effectively preserves both the sequential characteristics and feature relationships in the time-series data, making it suitable for tasks like stock price prediction where both realism and time-order are critical.

Data Set

Historical stock and index prices were obtained from Yahoo Finance, while the U.S. Dollar Index was retrieved from the Federal Reserve Economic Data (FRED) repository. Financial news articles were scraped from SeekingAlpha. Our primary target variable is the daily closing price of Apple Inc. From these closing prices, we derive all statistical features.

V. EXPERIMENT RESULT

Training Model

This study aims to forecast stock closing prices for the next three days using data from the previous 30 days. To train the prediction model, we incorporate not only historical closing prices but also 36 additional features that could influence price movements. The dataset is divided into 70% for training (1,726 records) and 30% for testing (739 records). The evaluation process includes two scenarios: one accounting for an unexpected event and one without such an event.

Experimental and results

RMSE

We assessed the performance of each model using the Root Mean Square Error (RMSE), calculated as:

$$RMSE = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \hat{x_i})^2}{N}}$$

where N is the number of data points, x_i is the actual stock price, and $\hat{x_i}$ is the predicted price. To evaluate the models developed in this project, we compared their RMSE values on test data, both including and excluding the year 2020

1. LSTM

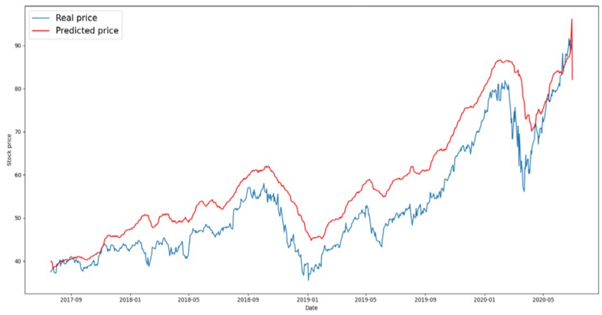

Fig 10.1

As shown in Fig. 10.1, the LSTM model's forecast including the year 2020 yields an RMSE of 6.60. In the graph, the blue line represents the actual stock price In contrast, Fig.10.2 illustrates the model's performance without including 2020, resulting in a significantly higher RMSE of 9.42.. In our LSTM model, the first layer employs a Bidirectional LSTM. We used the Adam optimizer with a learning rate of 0.001, a batch size of 64, and trained the model for 50 epochs on the stock price dataset

2. TimedGAN Model

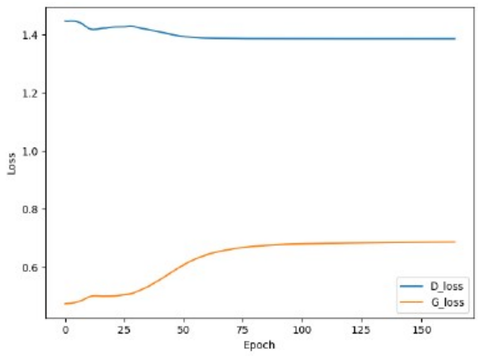

Figure 10.3 presents the loss curves of the *TimedGAN* model—where the blue line represents the discriminator's loss and the orange line shows the generator's loss. Initially, the discriminator's loss is higher than the generator's, but both gradually stabilize and flatten out as training progresses.

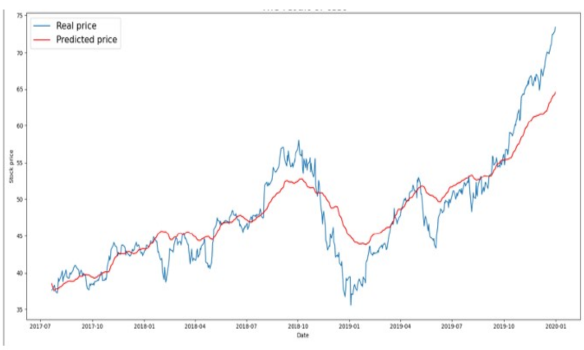

Fig.10.4

Fig 10.4 displays the model's predictions excluding data from 2020, where the RMSE drops to 3.09. This suggests that in the absence of unexpected events, basic TimedGAN model outperforms both baseline models in forecasting accuracy.

VI. Conclusion

We presented a TimeGAN framework employing a GRU-based generator and CNN-based discriminator for stock prediction. Experimental results show that both TimeGAN and WGAN-GP outperform LSTM and GRU baselines, with WGAN-GP demonstrating greater resilience during market disruptions and TimeGAN yielding superior accuracy under normal conditions. Despite these advantages, the integration of GANs with recurrent networks introduces training instability and heightened sensitivity to hyperparameter settings, which can impede model convergence and degrade predictive reliability if not carefully tuned. Future work will explore automated hyperparameter optimization and stability-enhancing regularization techniques.

References

[1] Q. M. Ilyas, K. Iqbal, S. Ijaz, A. Mehmood, and S. Bhatia, "A Hybrid Model to Predict Stock Closing Price Using Novel Features and a Fully Modified Hodrick–Prescott Filter," in Proc. Int. Conf. on Identification, Information and Knowledge in the Internet of Things (IIKI), 2018.

[2] K. Zhang, G. Zhong, J. Dong, S. Wang, and Y. Wang, "Stock Market Prediction Based on Generative Adversarial Network," Procedia Computer Science, vol. 147, pp. 400–406, 2019.

[3] X. Zhou, Z. Pan, G. Hu, S. Tang, and C. Zhao, "Stock Market Prediction on High-Frequency Data Using Generative Adversarial Nets," Mathematical Problems in Engineering, vol. 2018, Art. no. 4907423, 2018, doi: 10.1155/2018/4907423.

[4] Q. Bi, H. Yan, C. Chen, and Q. Su, "An Integrated Machine Learning Framework for Stock Price Prediction," in Proc. 26th China Conference on Information Retrieval (CCIR), Xi'an, China, Aug. 14–16, 2020.

[5] K. N. Rao, B. R. Gandhi, M. V. Rao, S. Javvadi, S. S. Vellela, and S. K. Basha, "Prediction and Classification of Alzheimer's Disease using Machine Learning Techniques in 3D MR Images," in Proc. Int. Conf. on Sustainable Computing and Smart Systems (ICSCSS), 2023.

[6] K. N. Rao, B. R. Gandhi, M. V. Rao, S. Javvadi, S. S. Vellela, and S. K. Basha, "Stock Values Predictions Using Deep Learning Based Hybrid Models," in Proc. Int. Conf. on Sustainable Computing and Smart Systems (ICSCSS), 2023.

[7] J. Yoon and D. Jarrett, "Time-series Generative Adversarial Networks," NeurIPS Time Series Workshop, 2019 https://papers.nips.cc/paper_files/paper/2019/file/c9efe5f2b76c6892ff49d8f9b31c0af1-Paper.pdf

[8] L. Mozaffari and J. Zhang, "Predicting Stock Prices: Strategies of Ensemble Learning with Transformer, ARIMA, and Linear Regression Models," SSRN, 2024 https://ssrn.com/abstract=4706455

[9] K. Uehara, "Deep Learning for Stock Market Prediction Using Sentiment and Technical Analysis," IEICE Transactions on Information and Systems, vol. E101.D, no. 4, pp. 1065–1072, Apr. 2018, doi: 10.1587/transinf.2016IIP0016.

[10] K. Uehara, "Stock Price Prediction by Deep Neural Generative Model of News Articles," in Proc. of the International Conference on Advances in Computing and Communication Engineering (ICACCE), 2022.

[11] Y. Li and Y. Pan, "A novel ensemble deep learning model for stock prediction based on stock prices and news," 2021.

[12] L. Mozaffari and J. Zhang, "Predicting stock prices: Strategies of ensemble learning with transformer, ARIMA, and linear regression models," 2024.

[13] "Deep learning for stock market prediction using sentiment and technical analysis," 2024.

[14] "Predicting stock market using natural language processing," 2022.

[15] K. Uehara, "Stock price prediction by deep neural generative model of news articles," 2022.

[16] S. Asgarian, R. Ghasemi, and S. Momtazi, "Generative adversarial network for sentiment-based stock prediction," *Concurrency and Computation: Practice and Experience*, vol. 35, no. 2, p. e7467, 2023. https://doi.org/10.1002/cpe.7467

[17]M. Staffini, "Stock price forecasting by a deep convolutional generative adversarial network," *Neural Computing and Applications*, vol. 34, pp. 1–13, 2022 https://doi.org/10.1007/s00521-022-07019-3.

[18] C. Yinka-Banjo, M. Akinyemi, and B. Er-rabbany, "Stock market prediction using a hybrid of deep learning models," *International Journal of Financial Studies, Economics and Management*, vol. 2, no. 2, pp. 1–16, 2023. https://doi.org/10.61549/ijfsem.v2i2.111.

[19] X. Wang, Y. Wang, B. Weng, and A. Vinel, "Stock2Vec: A hybrid deep learning framework for stock market prediction with representation learning and temporal convolutional network," *IEEE Access*, vol. 8, pp. 108–116, 2020 https://doi.org/10.1109/ACCESS.2020.2965436

[20] M. Nabipour, P. Nayyeri, H. Jabani, and A. Mosavi, "Deep learning for stock market prediction," *Applied Sciences*, vol. 10, no. 3, p. 850, 2020. https://doi.org/10.3390/app10030850