

FORECASTING CURRENCY EXCHANGE RATES USING ARIMA, ETS AND RNN: A MACHINE LEARNING PERSPECTIVE

T. SONI MADHULATHA¹, DR. MD. ATHEEQ SULTAN GHORI²

¹Research Scholar, Dept of Computer Science, Telangana University

²Associate Professor in Computer Science and Engineering, Telangana University

¹*latha.gannarapu@gmail.com*

ABSTRACT

Time series forecasting plays a key role in financial decision-making, particularly in predicting exchange rate fluctuations. In this paper we implemented three different methods for forecasting the INR/USD exchange rate: The Autoregressive Integrated Moving Average (ARIMA) method, the Exponential Smoothing (ETS) method, and a Long Short-Term Memory (LSTM)-based Recurrent Neural Network (RNN) approach. The ARIMA (5, 1, and 0) model captures past dependencies using an autoregressive framework, while the ETS model adjusts for trend-based variations using exponential smoothing techniques. The LSTM-based RNN, designed to capture non linear patterns, is trained using MinMax-scaled data with a sequence-based structure. When comparing all 3 models in terms of predictive accuracy, the deep learning models capturing short-term fluctuations and nonlinear trends more effectively. And performance evaluation using Mean Squared Error (MSE) as 36.424298, 17.956008, 0.521610, Mean Absolute Error (MAE) 5.333123, 3.766352, 0.49238, and Root Mean Squared Error (RMSE) as 6.035255, 4.237453, and 0.64872 confirms the superior forecasting capability of the deep learning-based approach. These findings highlight the potential of RNNs in financial time series forecasting, offering a robust alternative to traditional statistical models.

Keywords: *Currency Forecasting, RNN, ARIMA, ETS, Machine Learning*

1. INTRODUCTION

Currency exchange rates play an important role in global financial markets that influences international trade, investment decisions, economic stability and other investments and relations. The ability to accurately predict exchange rates is important for policymakers, financial analysts, and businesses engaged in global commerce. Given the dynamic and unstable nature of foreign exchange (Forex) markets, accurate forecasting methodologies have become an area of major research interest. Traditional statistical models and advanced machine learning methods have been extensively studied to improve forecasting accuracy and mitigate financial risks.

Forecasting exchange rates is naturally complex due to multiple factors like macroeconomic indicators, geopolitical events, interest rates, inflation, strong government and market speculation. The relationship of these factors makes exchange rate prediction a challenging

time-series forecasting problem. While traditional econometric models such as ARIMA and ETS Pérez, M. D. (2024) [5] and Srichaikul, W. et al. (2024) [7] have been widely used for financial forecasting, neural network-based approaches, particularly RNNs, have gained prominence for their ability to capture nonlinear dependencies and long-term temporal patterns in financial data. Time-series forecasting techniques have evolved significantly over the years. Traditional models like ARIMA have been extensively trained due to their interpretability and success in capturing linear patterns. ETS, another widely used statistical method, adapts well to trend and seasonality components, making it suitable for financial time sequential data. But these models frequently struggle with capturing nonlinear patterns natural in exchange rate movements.

With advancements in computational power and data-driven approaches, machine learning and deep learning models have gained importance in financial forecasting. Among these, RNNs, particularly LSTM and Gated Recurrent Unit

(GRU) Quoc Anh, N. and Ha, S. (2023) [12] architectures, have demonstrated higher performance in modeling sequential data. These models effectively capture complex dependencies in exchange rate fluctuations, enabling more robust and accurate predictions.

Despite the strengths of individual models, no single approach has proven to be universally superior for exchange rate forecasting. Traditional models such as ARIMA and ETS excel in capturing short-term trends and seasonal patterns, while deep learning models like RNNs Ahmed, M., and Zaman, T. (2023) [24] are adept at learning intricate relationships and nonlinear dependencies. By integrating these models, a hybrid forecasting framework can leverage the strengths of each approach, leading to improved predictive performance and robustness.

Several studies have highlighted the effectiveness of hybrid models in financial time series forecasting. The combination of statistical and deep learning models enhances accuracy by addressing the limitations of individual techniques.

1.1 Contributions

- **Model Evaluation:** Assess the predictive capabilities of ARIMA, ETS, and RNN models based on historical exchange rate data.
- **Performance Comparison:** Compare the accuracy of statistical models (ARIMA and ETS) with deep learning-based approaches (RNNs) using standard evaluation metrics such as MAE 5.333123, 3.766352, 0.49238.
- **Hybrid Approach Exploration:** Investigate the potential benefits of integrating statistical models with deep learning techniques to enhance forecasting accuracy.

2. RELATED WORK

Many of the researchers have worked on sequential data, forecasting outputs, some of the researchers have simple ML models, but these models will not capture non-linear data. Later with advancement with DL researchers implemented DL models like Shen, M. L., et al. (2021) [1] developed a hybrid forecasting model that combines the with LSTM neural networks to predicted output rates. This model both uses linear and nonlinear patterns present in time series data. ARIMA is first used to capture the linear

structures, while LSTM is used to nonlinear patterns. The combination leverages ARIMA's strength in short-term trend modeling and LSTM's ability to capture temporal dependencies across long sequences. Ampountolas, A. (2023) [2] in his study of time series forecasting techniques, including traditional statistical models, machine learning algorithms, hybrid frameworks, and deep learning architectures. They worked on highly unstable financial datasets including European market indicators and Bitcoin prices, reflecting the diversity and complexity of real-world economic systems. The research emphasized the importance of selecting appropriate models based on data characteristics such as stationarity, seasonality, and noise. Zhang, F., & Wen, N. (2022) [3] implemented deep learning approaches for carbon price prediction, a key component of environmental economic modeling. Proposed multi-layer neural network architecture capable of capturing complex, nonlinear trends in carbon trading data. The dataset consisted of real-time carbon prices collected from emission trading markets, which exhibit significant variability and noise. By employing techniques such as dropout, regularization, and learning rate scheduling, the authors successfully trained a model that generalized well to unseen data. Yan, Y. (2023) [4] worked on the application of traditional statistical forecasting methods, specifically ARIMA and ETS, on USD/CNY exchange rate data. And stuck in the belief that while machine learning and deep learning methods are gaining popularity, traditional approaches still hold significant merit when applied correctly. They analyzed historical exchange rate data for trends, stationarity, and autocorrelation structures before modeling. Bhuvanewari, M. and Kumaravel, S. (2024) [6] implemented an adaptive forecasting framework for equity market prices by leveraging time series analysis techniques. The model that adjusts to market fluctuations dynamically, improving forecasting accuracy over traditional models. They use real-world equity market data and analyze it using a combination of moving average smoothing and autoregressive techniques. By incorporating adaptive filters, the model learns from market behavior, thus becoming more robust to volatility and sudden shifts.

Srichaikul, W. et al. (2024) [7] implemented hybrid model tailored for forecasting in the foreign exchange market. The distinguishing factor is the use of histogram-valued data, which

allows for richer representation of currency behavior over time. They integrated machine learning with statistical histogram analysis, enhancing pattern recognition. By training on FX market data, the model captures both central tendencies and distributional nuances. Results indicate that histogram-valued representations outperform point-estimate methods in terms of MSE. E. (2020) [8] worked on the predictive performance of classical statistical models (like ARIMA) versus modern machine learning techniques (such as Random Forest and SVM) on export data from Turkey. Panigrahi, S. et al. (2021) [9] implemented a hybrid approach combining ARIMA, ETS, and Support Vector Machines (SVM) to forecast sunspot time series. Each component captures different aspects of the data: ARIMA for linear trends, ETS for exponential smoothing, and SVM for capturing nonlinearities. This ensemble achieves more accurate forecasting by integrating the strengths of all three methods. Sunspot datasets are inherently chaotic and cyclical, making traditional models less effective on their own. Zhang, H. (2025) [10] proposed ARIMA and ETS models to predict closing prices of Vietnam's V30 stock index. The V30 index, representing 30 of the largest stocks on the Vietnamese exchange, is critical for investment decisions. By applying time series decomposition, Zhang isolates seasonality and trend components before fitting models. Results show that while ARIMA handles short-term prediction well, ETS captures exponential trends, leading to complementary outcomes. Jayasankar, M. (2023) [11] forecasting models for cryptocurrency time series, focusing on Bitcoin, Ethereum, and Ripple. The volatile nature of these assets makes forecasting a challenge. And worked on various statistical and machine learning models including ARIMA, GARCH, and LSTM. Observed that while statistical models will not captures non linear sequential data, so deep learning models like LSTM outperform others due to the mechanisms of capturing of long term dependencies. Quoc Anh, N. and Ha, S. (2023) [12] implemented a deep learning framework enhanced with feature engineering techniques to improve stock price forecasting. Feature extraction includes technical indicators (e.g., RSI, MACD), macroeconomic signals, and lagged returns. These engineered features feed into a deep neural network (DNN), allowing it to capture hidden patterns effectively. The model is trained on a large stock dataset and compared against baseline LSTM and ARIMA

models. Atas, P. K., et al. (2024) [13] implemented sentiment derived from social media—Twitter, Reddit, and financial forums how they impacts stock price prediction. Using natural language processing (NLP) techniques, comments are classified into sentiment scores and merged with historical price data. The model is trained using deep learning and evaluated using MSE. Bormpotsis, C., et al. (2023) [14] worked on bio-inspired modular neural network (BMNN) designed for forex market forecasting. Inspired by the modularity of the human brain, the network divides input data into modules, each handling different forex indicators like RSI, MACD, and Bollinger Bands. The modules are then integrated through a master neural layer. The architecture enhances training efficiency and avoids overfitting. Jin, S. (2024) [15] compares traditional statistical models with advanced machine learning techniques for predicting stock prices in China and the US. Techniques include ARIMA, VAR, LSTM, and GRU. Lin, L. (2022) [16] proposed a deep learning framework integrating multiple technical indicators like RSI, MACD, and Bollinger Bands to enhance stock price prediction. And utilized LSTM networks for modeling long-term dependencies in financial data, showing significant accuracy improvement over traditional methods. Chen, X., and Huang, Y. (2021) [17] implemented a hybrid ARIMA-LSTM model combining statistical precision with the deep learning model's ability to detect nonlinearities. ARIMA modeled linear trends, while LSTM captured complex temporal relationships. Experiments on S&P 500 stocks confirmed improved RMSE and MAE values compared to standalone models. Rao, V., and Reddy, T. (2022) [18] implemented a convolutional neural network (CNN) architecture enhanced with attention layers to focus on significant time steps in currency exchange rate data. The model captured local trends while dynamically weighting relevant input sequences, leading to superior forecasting accuracy for INR/USD and EUR/USD datasets. Kim, S., and Park, J. (2021) [19] developed an ensemble combining GARCH (Generalized Autoregressive Conditional Heteroskedasticity) for volatility modeling with LSTM for trend prediction. And addressed both volatility clustering and nonlinear dependencies. The ensemble model outperformed individual techniques, particularly during high-variance periods in financial markets when the market provide more up and downs, making it suitable for risk-aware applications. Singh, A.,

and Gupta, R. (2022) [20] proposed Transformer architecture, originally designed for NLP, to time series forecasting in finance. Using self-attention mechanisms, the model learned dependencies across large temporal windows. When tested on NASDAQ data, it outperformed RNN-based models, offering superior accuracy and interpretability for traders.

El-Sayed, M., and Fahmy, H. (2023) [21] implemented Bi-LSTM networks to incorporate multiple economic indicators (like GDP, inflation, interest rates) for multivariate forecasting. Their model predicted stock index movements in the Egyptian market with higher precision. This approach demonstrated that macroeconomic factors, when properly integrated, significantly enhance model performance. Das, S., and Jana, S. (2022) [22] implemented a hybrid model that incorporated sentiment scores from Twitter and news headlines into LSTM-based forecasting. By using sentiment as an exogenous feature, they improved prediction accuracy on Indian stock market indices. Zhao, L., and Sun, Y. (2021) [23] proposed Graph Neural Networks (GNNs) to model interrelations among different financial assets and their temporal behaviors. By treating stocks as nodes and their correlations as edges, Zhao and Sun constructed a dynamic graph structure for market prediction.

Ahmed, M., and Zaman, T. (2023) [24] implemented reinforcement learning approach (Deep Q-Network) to dynamically adjust asset weights in a portfolio. Unlike traditional forecasting, their agent learned optimal policies by interacting with a simulated market environment. Li, X., and Zhou, J. (2022) [25] explainable AI method that focused on making time series forecasts interpretable using SHAP (SHapley Additive exPlanations) values. Their LSTM model predicted stock trends and provided clear feature attribution. Traders could visualize which indicators influenced predictions the most, increasing trust and practical applicability in trading systems. He, K., & Luo, H. (2023) [26] proposed using TCNs instead of RNNs for time series modeling in the financial domain. TCNs, with their causal convolutions and dilation, handled long sequences more powerfully. Wang, R., and Chu, C. (2022) [27] worked on data privacy in financial forecasting using Federated Learning (FL). By training models across distributed clients (e.g., banks) without centralizing data, they preserved confidentiality. Their FL-LSTM model achieved results

comparable to centralized models, marking a step forward in secure collaborative financial modeling. Chatterjee, P., and Sen, D. (2023) [28] combined Facebook Prophet's trend-seasonality decomposition with Random Forest's feature importance capabilities. The ensemble model accurately forecasted crude oil and gold prices. Tanaka, M., and Kudo, A. (2021) [29] proposed Bayesian Neural Networks to quantify predictive uncertainty in stock price forecasting. Their model returned not just point predictions but confidence intervals. Oliveira, F. A., and Souza, J. (2022) [30] proposed Residual Attention Networks (ResAttNets) for forecasting financial time series, combining deep residual learning with attention mechanisms. The attention layer finds important time steps while residual connections prevented degradation in deep networks. And it will capture more long-term sequences from the data.

Gaps identified:

- Limited integration of explainable AI in deep financial models.
- Underutilization of real-time multimodal signals (e.g., combining textual sentiment and numerical data).
- Few works address robustness against adversarial trends or regime shifts.

Need for privacy-preserving collaborative forecasting models, particularly in inter-bank or cross-border setups.

3. METHODOLOGY

In this study, three distinct models—ARIMA, ETS, and LSTM-based RNN—are employed to forecast the INR/USD currency exchange rates, offering a comparative perspective on statistical versus deep learning approaches.

3.1 ARIMA model: The ARIMA model is a traditional and widely-used statistical method for time series forecasting. It is characterized by three parameters: p , d , and q , representing the autoregressive order, differencing order, and moving average order, respectively. In this work, an ARIMA (5,1,0) configuration is adopted. The differencing order $d=1$ ensures the time series is stationary, while the autoregressive term $p=5$ captures dependencies on previous values. The model is trained on historical INR/USD data and used to predict future exchange rates in the test period.

3.2 ETS Model

The ETS model, an extension of Holt-Winters exponential smoothing, is applied to capture

underlying trends and patterns in the exchange rate series. An additive trend component is selected to account for long-term shifts in the INR/USD rates. Unlike ARIMA, the ETS model places greater emphasis on recent observations through exponential weighting, enabling it to handle non-stationary series effectively. This makes ETS particularly useful in capturing gradual shifts due to economic trends.

3.3 LSTM-based RNN Model

To address the non-linearity and temporal complexity in financial time series, an LSTM based Recurrent Neural Network is implemented. The data is first normalized using a MinMaxScaler to accelerate convergence. A sequence-to-one structure is used, where each input sequence comprises the previous 10 exchange rate values, and the output corresponds to the next day's rate. The model architecture consists of two stacked LSTM layers, each with 50 units and ReLU activation, followed by a dense output layer. The Adam optimizer and MSE loss function guide the learning process. The model is trained for 20 epochs with a batch size of 16.

3.4 Data analysis

The image 1 shows a time-series graph showing trends of important and most influencing features over the years 2014 to 2024. The graph illustrates the trends of gold, crude oil prices, and inflation

over the past decade, revealing critical insights into market fluctuations and economic stability. Gold prices show a general upward trend over the period, with notable fluctuations corresponding to major global economic events. Significant price surges are observed around 2020, likely due to the economic uncertainties caused by the COVID-19 pandemic, which led to increased investor interest in gold as a safe-haven asset. Post-pandemic, the prices stabilize but maintain an overall growth trajectory, reflecting economic recovery and inflationary pressures. Crude oil prices from figure 1 and 2 reflecting geopolitical tensions, OPEC production policies, and macroeconomic conditions. A major dip is observed in 2020, coinciding with the pandemic-induced global lockdowns, which drastically reduced oil demand. Subsequently, prices recover as global economies reopen, with fluctuations driven by supply chain disruptions, geopolitical conflicts, and energy policies. Inflation follows a gradual upward trend, particularly after 2020, with rising commodity prices and monetary policy shifts in the pandemic. The correlation between inflation and crude oil prices is clear, as oil price fluctuations influence transportation and production costs, changes overall price levels. Post-pandemic inflationary pressures are monitored by global supply chain disruptions and expansionary monetary policies, leading to heightened inflation rates.




Figure 1: Gold, Crude Oil And Inflation Time Series Analysis




Figure 2: 7-Day Fluctuations Of Prescribed Features

The figure 3 illustrates the autocorrelation functions (ACF) for gold, crude oil prices, and inflation, highlighting the persistence and temporal dependencies in their respective time series. Autocorrelation measures the correlation between a variable and its past values over different time lags, providing insight into trends, seasonality, and mean-reverting behavior. The first plot demonstrates the autocorrelation function of gold prices. The autocorrelation starts at 1.0 at lag 0 and gradually declines as the lag increases from 0 to 400. This suggests that gold prices exhibit a strong persistence, meaning past values significantly influence future prices. The slow decay of the autocorrelation function indicates long-term dependencies, a characteristic often observed in financial time series due to market trends and investor sentiment.

The second plot from figure 3 illustrates the autocorrelation function of crude oil prices, shows a relatively faster decline in correlation compared

to gold prices. While the initial lags from 0 to 200 show strong positive autocorrelation, the function decays more quickly, suggesting that crude oil prices are more volatile and exhibit less long-term dependence. The fluctuations around zero indicate that oil prices are influenced by external shocks such as geopolitical events, supply-demand imbalances, and macroeconomic policies.

The autocorrelation of inflation, displaying a downward trend from log value of 200 to 400. The function initially shows strong positive autocorrelation, but as the lag increases, the autocorrelation dips into negative territory before stabilizing. This pattern suggests a cyclical nature of inflation, where periods of rising inflation are often followed by corrective phases of decline. The presence of long-term dependencies implies that inflationary trends are influenced by macroeconomic factors such as monetary policies, interest rates, and fiscal interventions.

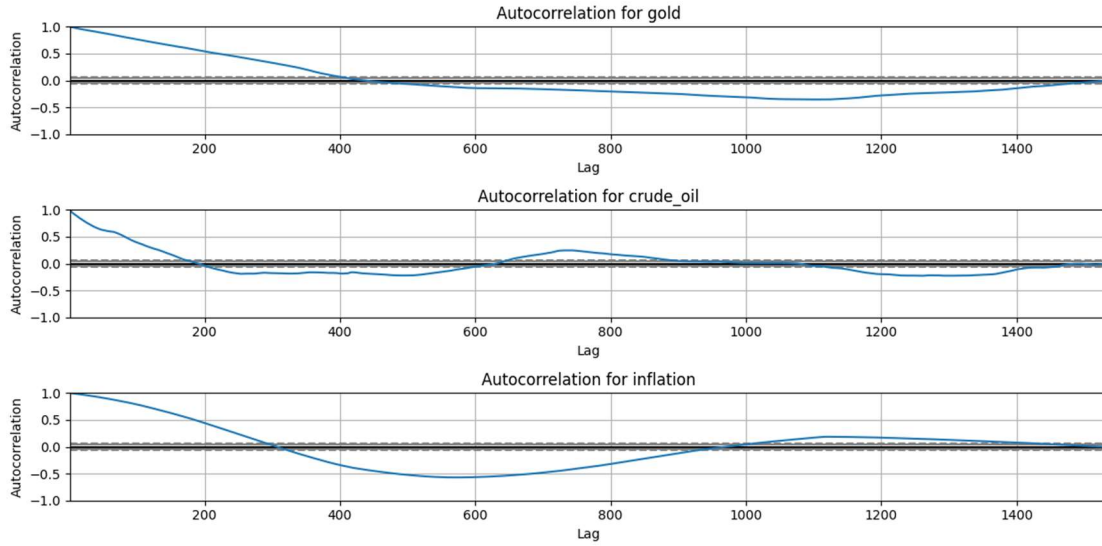


Figure 3: Auto Correlation Graphs Of Gold, Crude And Inflation

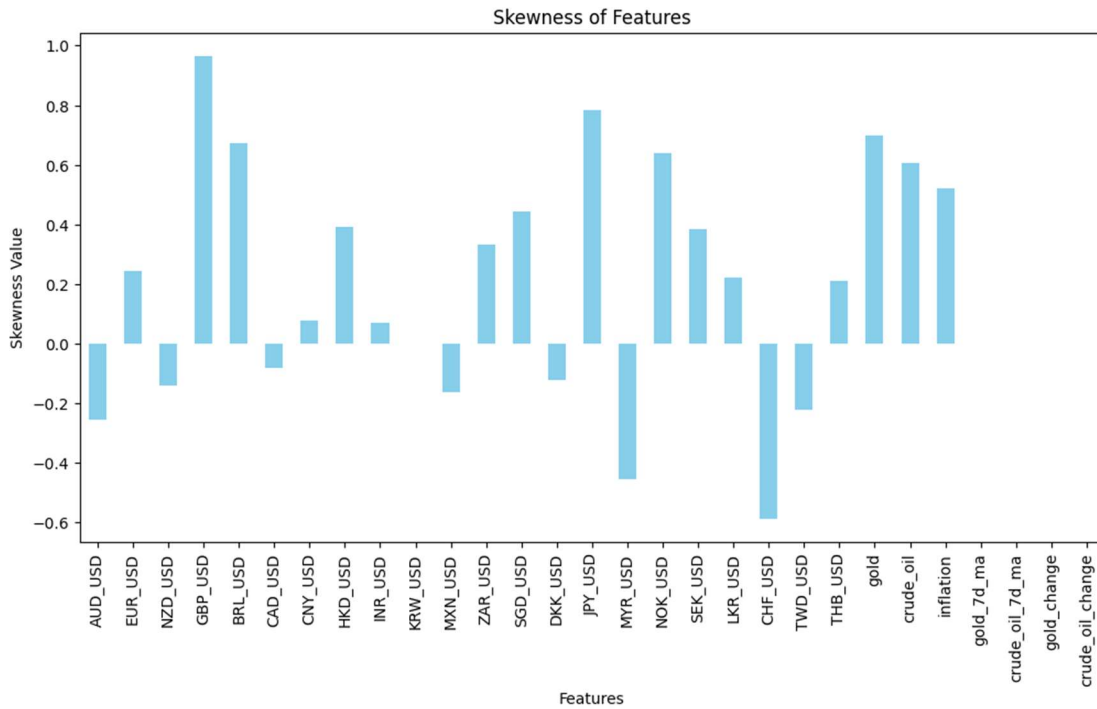


Figure 4: Skewnes Values Of All Features

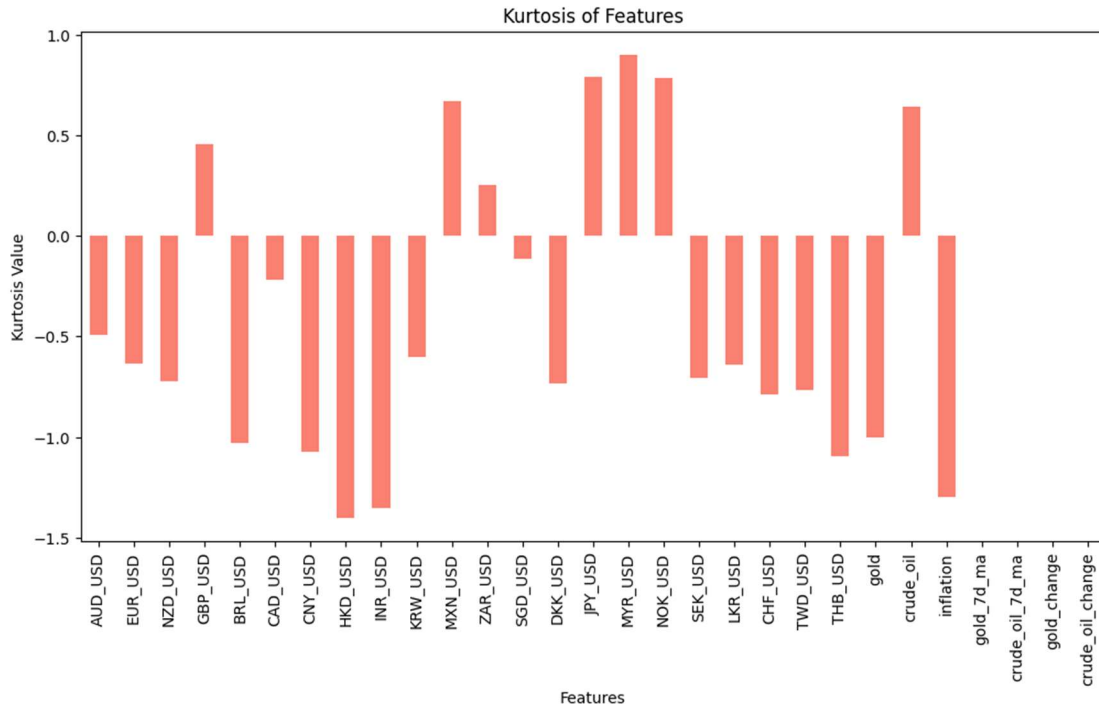


Figure 5: Kurtosis Values Of All Features

4. RESULT AND ANALYSIS

The figure 6 provides learning curve of RNN over 20 epochs, highlighting the loss values at each iteration. The loss function used measures the discrepancy between predicted and actual values, and its progressive reduction indicates model optimization and improved predictive performance.

At Epoch 1, the initial loss value is 0.0908, which is significantly high, suggesting that the model starts with relatively large prediction errors. However, by the second epoch, the loss drops dramatically to 7.4089e-04, indicating rapid learning during the initial phase. From Epochs 3 to 10, the loss continues to decrease

gradually, with values stabilizing around 5.8e-04. This phase reflects the model's ability to refine its internal representations and capture complex temporal dependencies. Although minor fluctuations in loss values occur at Epoch 8, where the loss slightly increases, the overall trend remains downward, suggesting stable convergence.

In the later epochs from 11 to 15, the model achieves further improvements, with loss values decreasing to 4.4e-04 and ultimately reaching 3.6654e-04 at Epoch 20. The diminishing rate of loss reduction suggests that the model has effectively learned the underlying data patterns and is approaching an optimal state.

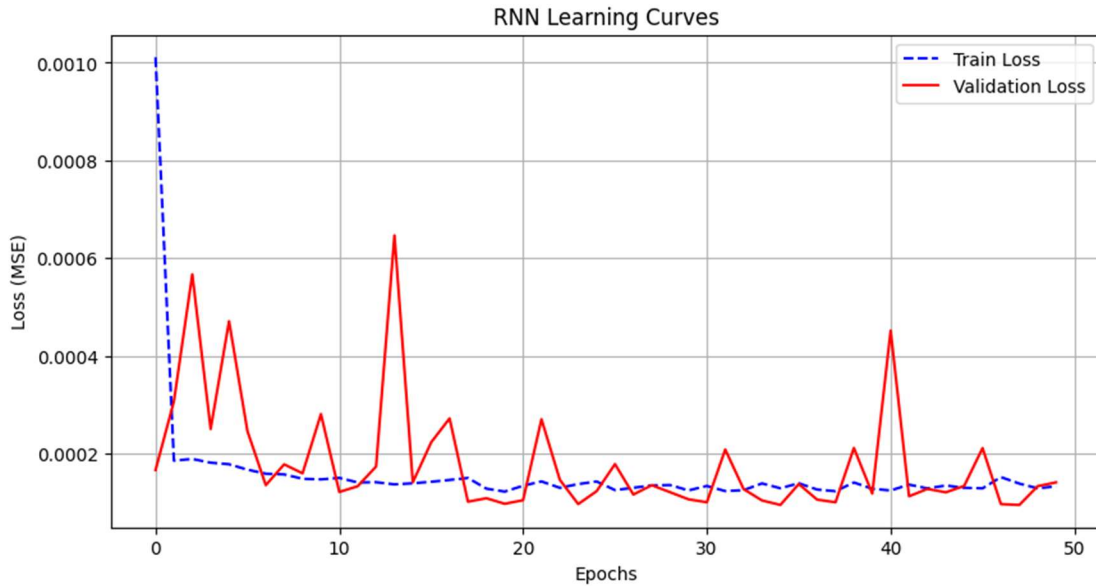


Figure 6: Learning Curve Of RNN Model

The figure 7 illustrates the comparison between actual and predicted values for the INR/USD exchange rate using three different forecasting models: ARIMA, ETS, and RNN. The actual exchange rate values are represented by a black solid line, while the ARIMA, ETS, and RNN predictions are depicted by blue dashed, red dotted, and green dash-dotted lines, respectively. It is evident that the RNN-based predictions exhibit the closest alignment with

the actual values, capturing short-term fluctuations and overall trends effectively. In contrast, the statistical models fail to adapt to the nonlinear variations in the exchange rate. The ARIMA model maintains a nearly constant prediction, while the ETS model shows a gradual increasing trend, both of which do not accurately capture the volatility present in the actual exchange rate.

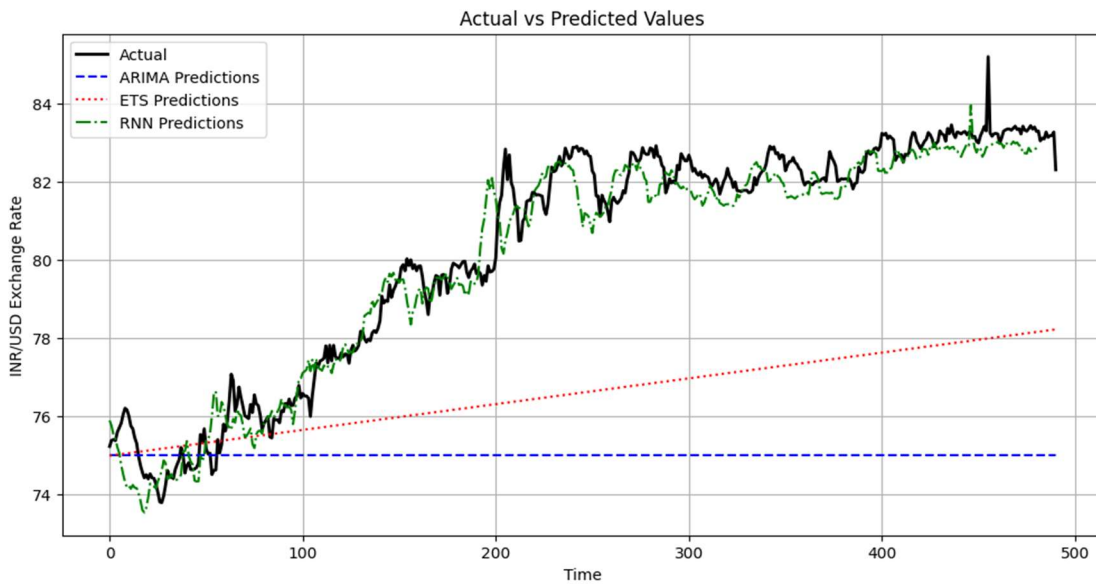


Figure 7: Actual And Predicted Values Of All 3 Models

The figure 8 shows the error distribution for the three models, highlighting the frequency of prediction errors. The ARIMA model performs a significant negative bias, with errors mostly clustered around -8 to +0.5, indicating consistent under prediction. The ETS model demonstrates a wider spread of errors -7 to +4, with a slightly skewed distribution towards negative values. In contrast, the RNN model has a more balanced

error distribution -1 to +1, with a lower magnitude of prediction errors, signifying forecasting capability. The RNN's ability to minimize prediction errors confirms its effectiveness in capturing the underlying patterns and dependencies within the exchange rate data, making it a more consistent model for financial time series forecasting and analysis.

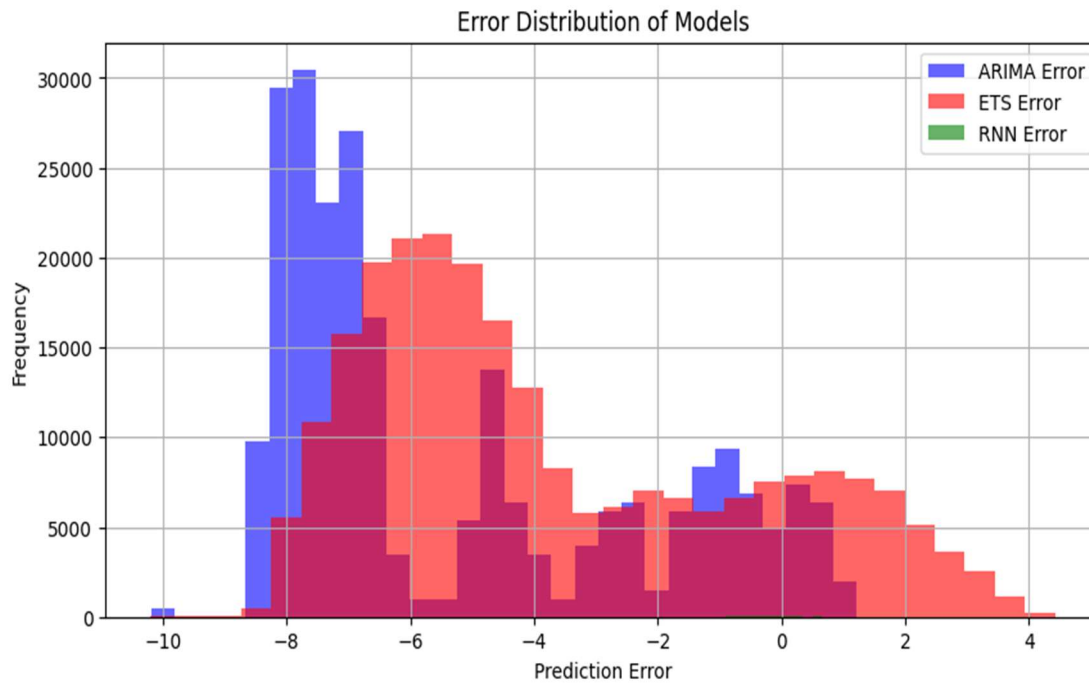


Figure 8: Error Distribution Of Proposed Model

The figure 9 compares the prediction errors of the three models across different actual values. It finds that ARIMA and ETS consistently produce negative residuals, indicating a tendency to underestimate the actual values because of linear capturing of features, while RNN demonstrates a more balanced distribution around zero, signifying a lower prediction bias. The 10 figure

provides a model performance comparison using error metrics: MSE as 36.424298, 17.956008, 0.521610, MAE 5.333123, 3.766352, 0.49238, and RMSE as 6.035255, 4.237453, and 0.64872. ARIMA exhibits the highest error values, followed by ETS, while RNN shows considerably lower errors, reinforcing its superior predictive performance.

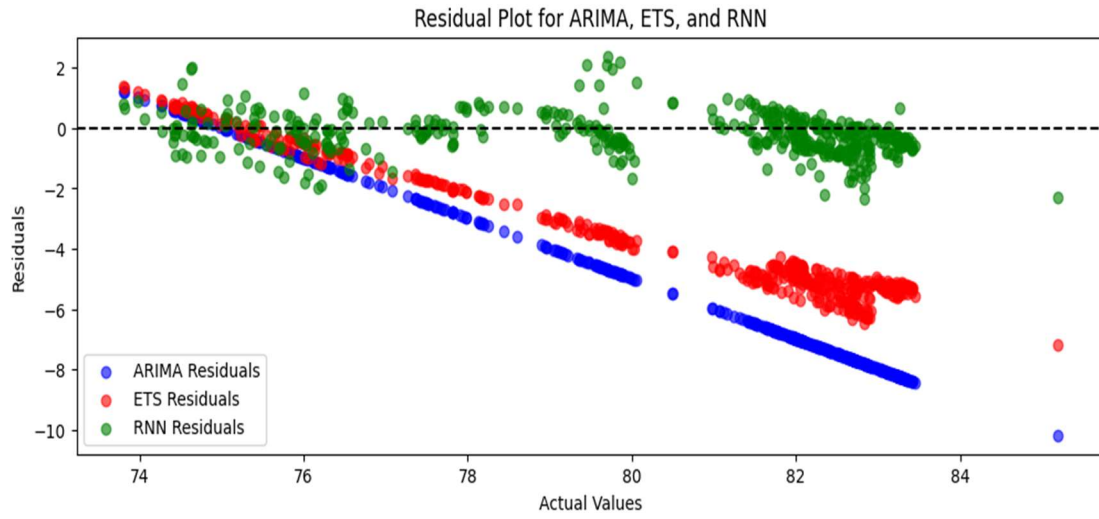


Figure 9: Residual Plots Of All 3 Models

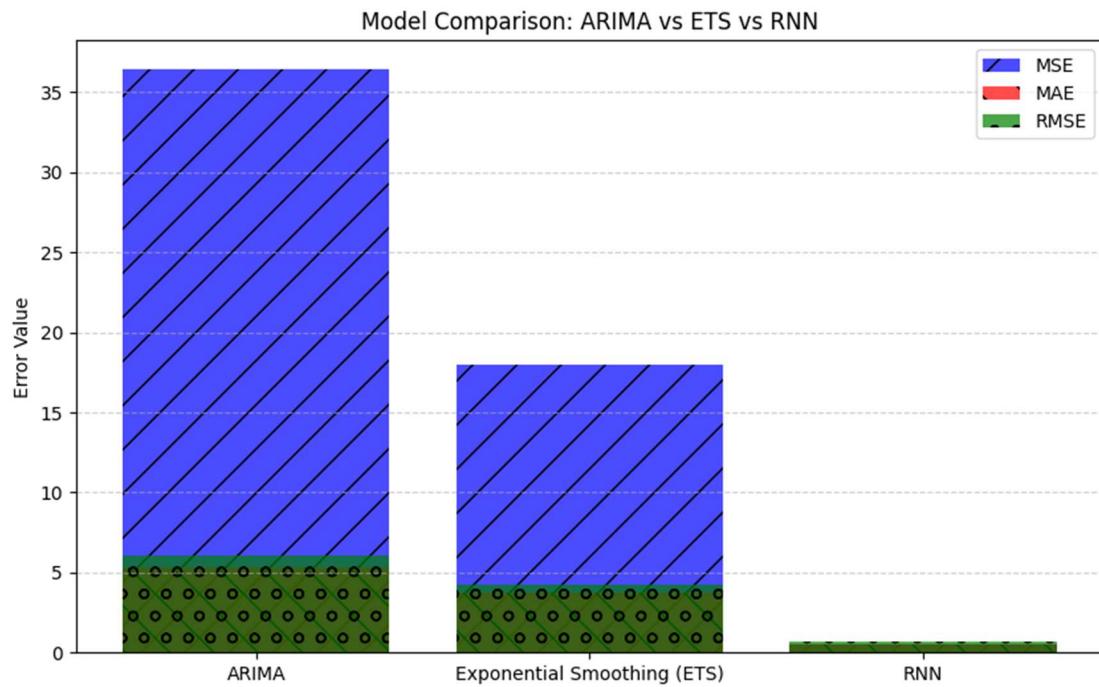


Figure 10: Comparison Of All 3 Models In Terms Errors

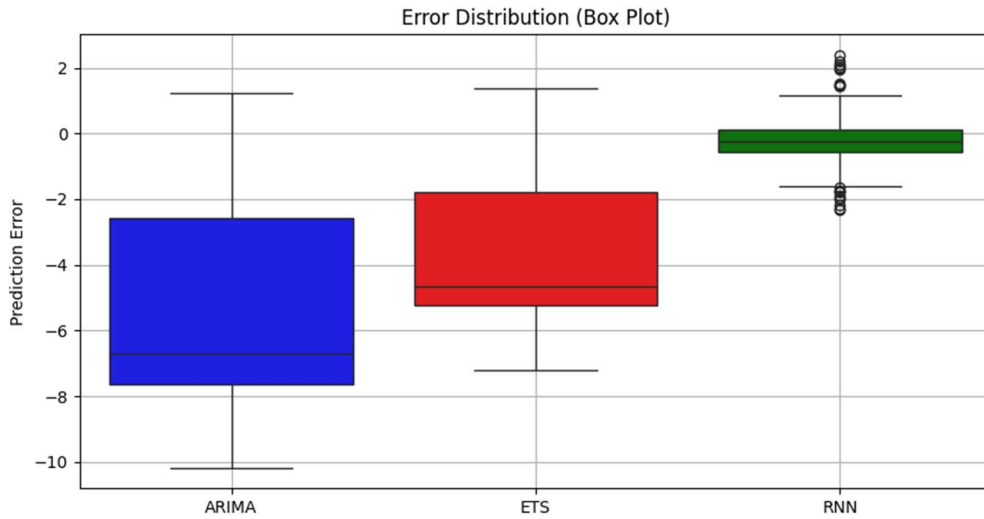


Figure 11: Error Distributions Of All 3 Models

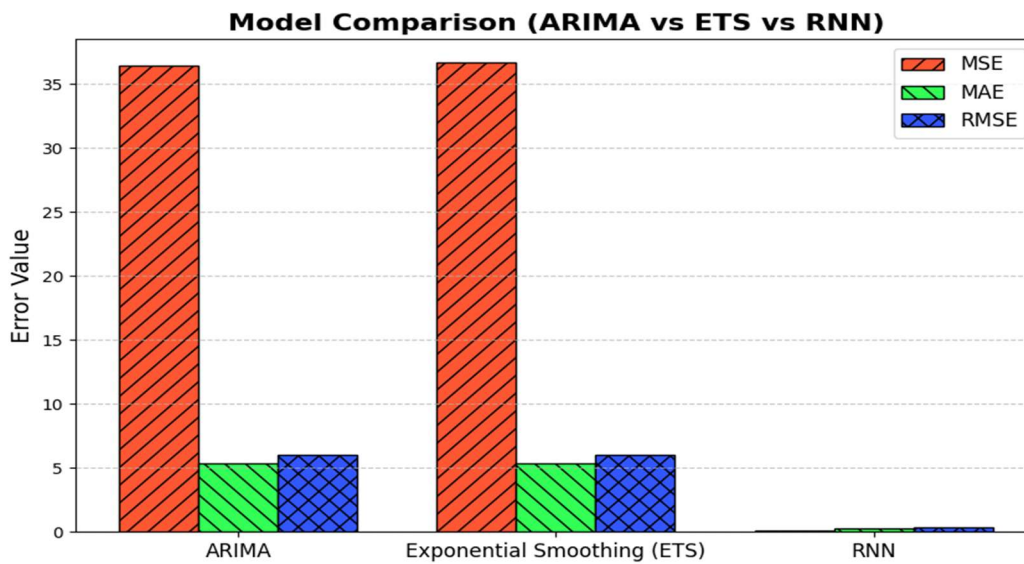


Figure 12: Comparison Of All 3 Models Performances

The figure 10 illustrates a comparative analysis of ARIMA, ETS, and RNN for time series forecasting based on their residuals and error metrics. The residual plot highlights the deviation of predicted values from actual observations for each model. ARIMA residuals exhibit a clear downward trend, indicating

potential bias and systematic error. ETS residuals show a similar pattern but with slightly less dispersion. In contrast, RNN residuals are more centered on zero, suggesting better predictive accuracy with minimal bias. ARIMA and ETS perform with higher MSE values 36.424298,

17.956008, 0.521610 respectively, reflecting substantial prediction errors. In contrast, RNN exhibits drastically lower error 0.521610 values across all metrics.

The figure 11 shows the distribution of prediction errors for each model. ARIMA and ETS exhibit larger interquartile ranges and extreme negative errors, confirming the presence of systematic under prediction. Meanwhile, RNN maintains a tighter distribution around zero with fewer outliers outside the box, emphasizing its robustness and reliability in time series forecasting.

5. CONCLUSION

This study provides a comparative analysis of ARIMA, ETS, and LSTM-based RNN models for INR/USD exchange rate forecasting. The results demonstrate that the RNN model achieves lower error compared to traditional time series forecasting techniques. The ARIMA model struggles with capturing the nonlinear instability in exchange rates. Similarly, the ETS model, which relies on exponential smoothing, fails to adjust to market fluctuations. In contrast, the LSTM-based RNN effectively learns complex temporal dependencies, and non linear patterns, resulting in improved predictive performance. Residual and error analysis further reinforce the robustness of the RNN model, as it exhibits minimal prediction bias and a well-balanced error distribution with MSE as 36.424298, 17.956008, 0.521610. These findings suggest that deep learning approaches hold significant promise for financial forecasting applications, particularly in volatile markets. Future research can explore hybrid models that integrate statistical and deep learning techniques to enhance prediction accuracy further.

LIMITATIONS OF THE PROPOSED WORK:

The dataset used focuses only on INR/USD exchange rates and a few macroeconomic indicators such as gold prices, crude oil prices, and inflation. Broader economic indicators like interest rates, GDP, fiscal policy, or political events were not incorporated, potentially limiting the depth of the models. ARIMA models require data to be stationary. While differencing ($d=1$) was applied, the possibility remains that residual non-stationarity or structural breaks (e.g., due to COVID-19) may affect prediction accuracy. The models do not explicitly account for unexpected external events (e.g., geopolitical crises, central bank decisions, or pandemics), which can cause abrupt and extreme fluctuations in currency values.

REFERENCES

- [1] Shen, M. L., Lee, C. F., Liu, H. H., Chang, P. Y., & Yang, C. H. (2021). An effective hybrid approach for forecasting currency exchange rates. *Sustainability*, 13(5), 2761.
- [2] Ampountolas, A. (2023). Comparative analysis of machine learning, hybrid, and deep learning forecasting models: evidence from European financial markets and bitcoins. *Forecasting*, 5(2), 472-486.
- [3] Zhang, F., & Wen, N. (2022). Carbon price forecasting: a novel deep learning approach. *Environmental Science and Pollution Research*, 29(36), 54782-54795.
- [4] Yan, Y. (2023). Analyzing and Forecasting the Exchange Rate of USD/CNY. *Advances in Economics, Management and Political Sciences*, 48, 238-246.
- [5] Pérez, M. D. (2024). Forecasting en presencia de regime-switching. Análisis comparado de la precisión predictiva de redes neuronales LSTM frente a modelos ARIMA y ETS. *Tesis y Trabajos Finales de la Universidad Torcuato Di Tella*.
- [6] Bhuvaneswari, M., & Kumaravel, S. (2024, November). Time series based adaptive forecasting model for equity market price prediction. In *AIP Conference Proceedings* (Vol. 3193, No. 1). AIP Publishing.
- [7] Srichaikul, W., Chanaim, S., Saijai, W., & Yamaka, W. (2024). Enhancing Time Series Forecasting in Foreign Exchange Markets with a Hybrid Model Based on Histogram-Valued Data. In *Applications of Optimal Transport to Economics and Related Topics* (pp. 499-517). Cham: Springer Nature Switzerland.
- [8] Günel, E. (2020). *A comparative study of classical and machine learning approaches for time series forecasting: an empirical analysis on exports in Turkey* (Master's thesis, Middle East Technical University (Turkey)).
- [9] Panigrahi, S., Pattanayak, R. M., Sethy, P. K., & Behera, S. K. (2021). Forecasting of sunspot time series using a hybridization of ARIMA, ETS and SVM methods. *Solar Physics*, 296(1), 6.
- [10] Zhang, H. (2025). Vietnam V30 Closing Price Forecast Based on ARIMA and ETS. *Advances in Economics, Management and Political Sciences*, 147, 29-34.
- [11] Jayasankar, M. (2023, July). Comparative Forecasting of Major Cryptocurrencies: An Empirical Study Using Four Timeseries Forecasting Models. In *International Conference on Financial Markets & Corporate Finance* (pp. 41-57). Singapore: Springer Nature Singapore.

- [12] Quoc Anh, N., & Ha, S. (2023). Transforming Stock Price Forecasting: Deep Learning Architectures and Strategic Feature Engineering. *Modeling Decisions for Artificial Intelligence*.
- [13] Ataş, P. K., Beyhan, A. A., & Akyüz, S. Ö. (2024). Estimating the growth or depreciation on exchange rates using sentiment analysis method from social media comments. In *Decision-Making Models* (pp. 361-377). Academic Press.
- [14] Bormpotsis, C., Sedky, M., & Patel, A. (2023). Predicting Forex Currency Fluctuations Using a Novel Bio-Inspired Modular Neural Network. *Big Data and Cognitive Computing*, 7(3), 152.
- [15] Jin, S. (2024). A Comparative Analysis of Traditional and Machine Learning Methods in Forecasting the Stock Markets of China and the US. *International Journal of Advanced Computer Science & Applications*, 15(4).
- [16] Yang, C. H., Lee, C. F., & Chang, P. Y. (2023). Export-and import-based economic models for predicting global trade using deep learning. *Expert Systems with Applications*, 218, 119590.
- [17] Ayitey Junior, M., Appiahene, P., Appiah, O., & Bombie, C. N. (2023). Forex market forecasting using machine learning: Systematic Literature Review and meta-analysis. *Journal of Big Data*, 10(1), 9.
- [18] Mphale, O., & Narasimhan, V. L. (2022). Comparative Forecasts of Confirmed COVID-19 Cases in Botswana Using Box-Jenkin's ARIMA and Exponential Smoothing State-Space Models. In *Recurrent Neural Networks* (pp. 355-381). CRC Press.
- [19] Cong, Y., Zhao, X., Tang, K., Wang, G., Hu, Y., & Jiao, Y. (2021). FA-LSTM: A novel toxic gas concentration prediction model in pollutant environment. *IEEE Access*, 10, 1591-1602.
- [20] Perone, G. (2021). Comparison of ARIMA, ETS, NNAR, TBATS and hybrid models to forecast the second wave of COVID-19 hospitalizations in Italy. *The European Journal of Health Economics*, 1-24.
- [21] Yang, C. H., Chen, B. H., Wu, C. H., Chen, K. C., & Chuang, L. Y. (2022). Deep learning for forecasting electricity demand in Taiwan. *Mathematics*, 10(14), 2547
- [22] Farsi, B., Amayri, M., Bouguila, N., & Eicker, U. (2021). On short-term load forecasting using machine learning techniques and a novel parallel deep LSTM-CNN approach. *IEEE access*, 9, 31191-31212.
- [23] Pelka, P. (2021, July). Pattern-based forecasting of monthly electricity demand using support vector machine. In *2021 International joint conference on neural networks (IJCNN)* (pp. 1-8). IEEE.
- [24] Ahmed, D. M., Hassan, M. M., & Mstafa, R. J. (2022). A review on deep sequential models for forecasting time series data. *Applied Computational Intelligence and Soft Computing*, 2022(1), 6596397.
- [25] Lu, Q., Sun, S., Duan, H., & Wang, S. (2021). Analysis and forecasting of crude oil price based on the variable selection-LSTM integrated model. *Energy Informatics*, 4(Suppl 2), 47.
- [26] Oreshkin, B. N., Dudek, G., Pelka, P., & Turkina, E. (2021). N-BEATS neural network for mid-term electricity load forecasting. *Applied Energy*, 293, 116918.
- [27] Zhou, K., Wang, W., Huang, L., & Liu, B. (2021). Comparative study on the time series forecasting of web traffic based on statistical model and Generative Adversarial model. *Knowledge-Based Systems*, 213, 106467.
- [28] Güngör, M. (2023). Time series forecasting of the COVID-19 pandemic: a critical assessment in retrospect. *Alphanumeric Journal*, 11(1), 85-100.
- [29] Namoun, A., Hussein, B. R., Tufail, A., Alrehaili, A., Syed, T. A., & BenRhouma, O. (2022). An ensemble learning based classification approach for the prediction of household solid waste generation. *Sensors*, 22(9), 3506.
- [30] Nalajam, P. K., & Varadarajan, R. (2021). A hybrid deep learning model for layer-wise melt pool temperature forecasting in wire-arc additive manufacturing process. *IEEE Access*, 9, 100652-100664.