Predicting Foreign Exchange Rates for an Oil-Rich Country using Energy-Based Indicators

Jade Soodoo
Department of Computing and Information Technology
The University of the West Indies
St. Augustine, Trinidad
jade.soodoo@my.uwi.edu

Patrick Hosein
Department of Electrical and Computer Engineering
The University of the West Indies
St. Augustine, Trinidad
patrick.hosein@uwi.edu

***Abstract*—Trinidad and Tobago is an energy-rich country and hence its revenue is highly dependent on oil and gas prices. Since its foreign income is primarily based on oil and gas sales, this heavy dependency significantly affects the exchange rates of various currencies with some more dependent on oil and gas prices than others. This study explores forecasting the monthly buying and selling rates of eight major currencies. Given the volatility of global energy markets, accurate forecasting is essential for economic planning and stability. The study incorporates energy-based predictors such as crude oil production, methanol exports and natural gas output, to improve forecast accuracy. A Random Forest Regressor identified the most influential indicators for each currency, followed by XGBoost, LSTM, and GRU models. Four ensemble models were created from a combination of LSTM, GRU, and XGBoost models. Rolling origin forecasting was used to replicate real-time scenarios and performance was evaluated using MAPE, MAE, and NRMSE metrics. Ensemble models performed significantly better than single models. Results confirmed that integrating energy indicators enhances predictive accuracy, offering valuable insights to central banks and policymakers managing exchange rates. We believe that this analysis will hold for similar energy-rich states.**

***Index Terms*—Exchange Rate Forecasting, Energy Commodities, Machine Learning Models, Rolling Origin Forecasting, Trinidad and Tobago Economy**

I. Introduction

Trinidad and Tobago's economy is highly dependent on the energy sector and international trade. As a small island developing state (SIDS), the country exports energy products and imports a large number of goods required for local consumption. Hence, exchange rate fluctuations have significant implications, affecting the prices of food, fuel, medicine, construction materials, and other essential commodities. Fluctuations in exchange rates also influence inflation, trade competitiveness, and consumer expenditure. While the Central Bank of Trinidad and Tobago (CBTT) manages the exchange rate within a managed float system, both global and domestic shocks continue to impact market dynamics.

As energy exports account for a significant percentage of the country's foreign exchange earnings, changes in production and export levels can have a significant impact on the value of the Trinidad and Tobago dollar (TTD).

978-1-6654-7808-3/25/$31.00 ©2025 IEEE

The ongoing volatility in international energy markets highlights the necessity for improved exchange rate forecasting techniques that will be sensitive to actual economic conditions in the world. Traditional forecasting models depend too much on linear relationships and smooth trends and as such, fail to do well in trying to model the intricate, nonlinear dynamics of an energy-based economy.

This study addresses these challenges by applying machine learning and ensemble models to forecast the monthly buying and selling rates of eight major currencies relevant to Trinidad and Tobago's trade and investment landscape: BBD, CAN, CHF, XCD, GBP, JPY, USD, and EURO. Monthly exchange rate data from January 1999 to December 2023 are combined with 16 energy-based indicators inclusive of crude oil production, natural gas output, and methanol exports, sourced from the Central Bank of Trinidad and Tobago's Data Centre.

To ensure rigorous evaluation, a rolling origin (time-series cross-validation) approach is employed, simulating real-world forecasting conditions. The random forest regression coefficient is first used to identify the two main predictors for each currency, both the buying and selling rate, followed by the development of XGBoost, LSTM and GRU models. Ensemble models were also developed by combining forecasts from multiple learners. Performance is assessed using Mean Absolute Percentage Error (MAPE), Mean Absolute Error (MAE), and Normalized Root Mean Square Error (NRMSE). Visualizations will also be provided to support comparative analysis.

This research is significant as it combines actual-world economic measures with sophisticated forecasting techniques. The findings are beneficial to policymakers, central banks, and financial institutions that seek to enhance foreign exchange management in energy-exporting economies.

II. Related Work

Exchange rate prediction is a long-standing issue in economics and finance due to the highly volatile and multicausal nature of currency markets. Influences range from macroeconomic trends and geopolitical events to commodity prices [1]. For energy-dependent economies such as Trinidad and Tobago, the prices and production levels of energy based commodities such as crude oil, natural gas, and petrochemical exports play a significant role in exchange rate movements [2].

Energy-exporting nations typically experience currency appreciation when global energy prices rise, while importing economies may face depreciation under the same conditions [3]. To capture this relationship, recent studies incorporate production metrics of energy commodities—like crude oil, natural gas, ammonia, methanol, and fertilizers—as explanatory variables in forecasting models [2].

The foundational work by Meese and Rogoff [1] famously showed that traditional structural models failed to outperform a naïve random walk in out-of-sample exchange rate prediction. This sparked ongoing research into more advanced modeling strategies. While early models such as ARIMA and SARIMA were widely used due to their mathematical clarity [4], they often failed in the presence of nonlinearities or structural breaks [5].

To transcend these shortcomings, consideration was shifted to machine learning (ML) and deep learning (DL) approaches. More particularly, XGBoost, Long Short-Term Memory (LSTM) networks, and Gated Recurrent Units (GRU) models have been shown to have great promise in modeling complex time-dependent patterns in exchange rate data.

The random forest regression (RFR) has emerged as a powerful tool for feature selection, capable of uncovering nonlinear dependencies and ranking the importance of the variable [6]. It has been widely used in economic forecasting to detect critical relationships between macroeconomic variables and target outcomes [7].

Gradient Boosting Machines (GBMs), particularly XGBoost, are also prominent in financial modeling. Unlike Random Forests, which train trees in parallel, GBMs build them sequentially to correct prior errors, leading to better accuracy with careful tuning [8]. XGBoost is especially adept at capturing sudden shocks or nonlinear jumps in data, as confirmed in studies comparing its performance to LSTM models [9].

LSTM networks are a type of recurrent neural network (RNN) specifically designed to capture long-term dependencies in time series data [10]. Hybrid LSTM models that combine with traditional statistical approaches like ARIMA have further improved performance by addressing both linear and nonlinear components [11].

GRUs are a simplification of LSTMs where the input and forget gates are combined into a single update gate [12]. Comparative studies have determined that GRUs can match or outperform LSTM performance in the majority of financial forecast use cases [10], [12].

Despite the strengths of an individual model, no single method dominates others in all environments [13]. This inconsistency has led to the rise of ensemble models, which combine the predictions of multiple algorithms to reduce error, increase robustness, and generalize better to new data [14].

These ensemble approaches not only improve accuracy, but also stabilize performance between currencies with varying sensitivities to energy-related indicators [13], [14].

III. Methodology

A. Data Collection and Preparation

Exchange rate data for eight currencies (BBD, CAN, CHF, XCD, GBP, JPY, USD, EURO) from January 1999 to December 2024 were obtained from the Central Bank of Trinidad and Tobago. Energy commodity data, including production and export volumes of crude oil, LNG, methanol, and urea, were merged on the "Date" column. Missing values were replaced using corresponding yearly averages.

B. Feature Selection

A Random Forest Regressor (RFR) was used to identify the most relevant energy indicators for each currency. For a Random Forest with T trees, the prediction $\hat{y}$ for feature vector x is:

$$\hat{y} = \frac{1}{T}\sum_{t=1}^{T} h_t(x) \tag{1}$$

where h_t is the prediction from tree t. Feature importance for feature j was computed as the average decrease in impurity brought by that feature across all trees.

C. Model Evaluation Framework

1) 80–20 Train-Test Split: As a benchmark, models were first evaluated using an 80% training and 20% testing split for four configurations per currency using the XGBoost algorithm:

1) No predictors
2) First selected predictor
3) Second selected predictor
4) Both predictors

The configuration that yielded the lowest mean absolute percentage error (MAPE) value was selected as the optimal setup for future forecasting.

2) Rolling Origin Forecasting: For the final evaluation, a rolling origin approach was used. For each month in 2024, models were trained on all historical data up to the previous month:

$$\hat{y}_{t+1} = f(y_1, y_2, \ldots, y_t) \tag{2}$$

This ensured no data leakage and realistic forecasting.

D. Forecasting Models

Three forecasting models were implemented: Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Extreme Gradient Boosting (XGBoost). These were selected because they are well-known for handling temporal dependencies and nonlinear relationships in time-series forecasting, particularly in financial and economic domains.

1) Extreme Gradient Boosting (XGBoost): XGBoost is a scalable tree boosting algorithm that builds an ensemble of decision trees sequentially to minimize a regularized objective function. It is particularly effective at capturing nonlinear interactions between predictors and handling heterogeneous data types. The model minimizes:

$$\mathcal{L}(\phi) = \sum_i l(\hat{y}_i, y_i) + \sum_k \Omega(f_k) \tag{3}$$

where

$$\Omega(f_k) = \gamma T + \frac{1}{2}\lambda\|w\|^2 \tag{4}$$

controls tree complexity and prevents overfitting. XGBoost was chosen as a benchmark due to its speed, interpretability via feature importance, and robustness to missing values.

2) Long Short-Term Memory (LSTM): LSTM networks are a variant of recurrent neural networks (RNNs) designed to capture long-term dependencies by mitigating the vanishing gradient problem. They maintain a memory cell C_t that stores contextual information across time steps, regulated by input (i_t), forget (f_t), and output (o_t) gates:

$$f_t = \sigma(W_f[h_{t-1}, x_t] + b_f) \tag{5}$$

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i) \tag{6}$$

$$\tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C) \tag{7}$$

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t \tag{8}$$

$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o) \tag{9}$$

$$h_t = o_t * \tanh(C_t) \tag{10}$$

This structure enables the model to capture both short-term fluctuations and persistent patterns in exchange rate movements.

3) Gated Recurrent Unit (GRU): GRUs are a simplified alternative to LSTMs that merge the forget and input gates into a single update gate z_t and use a reset gate r_t. This reduces the number of trainable parameters and often results in faster training:

$$z_t = \sigma(W_z[h_{t-1}, x_t]) \tag{11}$$

$$r_t = \sigma(W_r[h_{t-1}, x_t]) \tag{12}$$

$$\tilde{h}_t = \tanh(W_h[r_t * h_{t-1}, x_t]) \tag{13}$$

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t \tag{14}$$

GRUs are computationally efficient and particularly useful when training time is constrained while still preserving strong sequence modeling capabilities.

E. Ensemble Models

To improve accuracy and robustness, ensemble learning was employed by averaging forecasts from different models. The rationale is that individual models have different inductive biases and error patterns; combining them can reduce variance and capture more comprehensive dynamics. Four ensembles were tested:

- **LSTM + GRU:** Combines two sequence models with different gating mechanisms to balance overfitting tendencies and strengthen temporal pattern capture.
- **LSTM + XGBoost:** Leverages LSTM's ability to model sequential dependencies and XGBoost's strength in handling nonlinear relationships and structured features.
- **GRU + XGBoost:** Offers a balance between computational efficiency and representational power, suitable for rapid deployment.
- **LSTM + GRU + XGBoost:** The most comprehensive ensemble, combining deep sequential learning with tree-based structure learning to maximize predictive stability.

Final ensemble forecasts were computed as inverse-MAPE weighted averages of component model predictions, giving higher influence to more accurate models.

F. Performance Metrics

Model performance was assessed using three complementary metrics to capture different error characteristics.

1) Mean Absolute Percentage Error (MAPE): Measures average percentage deviation between forecasts and actual values, making it scale-independent:

$$\text{MAPE} = \frac{100\%}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \tag{15}$$

2) Mean Absolute Error (MAE): Calculates the mean magnitude of errors in original units, treating all errors equally:

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{16}$$

3) Normalized Root Mean Squared Error (NRMSE): Penalizes larger errors more heavily and normalizes by the mean of actual values for comparability across currencies:

$$\text{NRMSE} = \frac{\sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}}{\bar{y}} \tag{17}$$

where $\bar{y}$ is the mean observed value.

IV. Results

A. Exploratory Data Analysis

To gain initial insights into the distribution and variability of the energy commodity indicators, histograms with superimposed kernel density estimates (KDEs) were produced for all variables. These plots reveal skewness, multimodality, and potential outliers.

Figure 1 shows that most energy commodities are positively skewed, with some exhibiting multimodal patterns, reflecting periods of unusually high production or exports, structural changes, or seasonal effects. The long right tails and outliers indicate volatility that could impact model accuracy. In contrast, a few commodities, such as ammonia and fertilizer exports, have relatively symmetric distributions, suggesting greater stability.

The correlation heatmap highlights pairwise Pearson correlation coefficients among exchange rates and energy indicators, which guide the selection of predictors. Figure 2 reveals strong positive correlations within variable groups, such as between buying and selling rates of the same currency and between related energy indicators like crude oil production and exports. Moderate positive correlations between some exchange rates and energy variables, particularly crude oil and methanol, indicate commodity market influence on currency levels. In contrast, weak or negative correlations, such as between the JPY rate and certain energy exports, show that

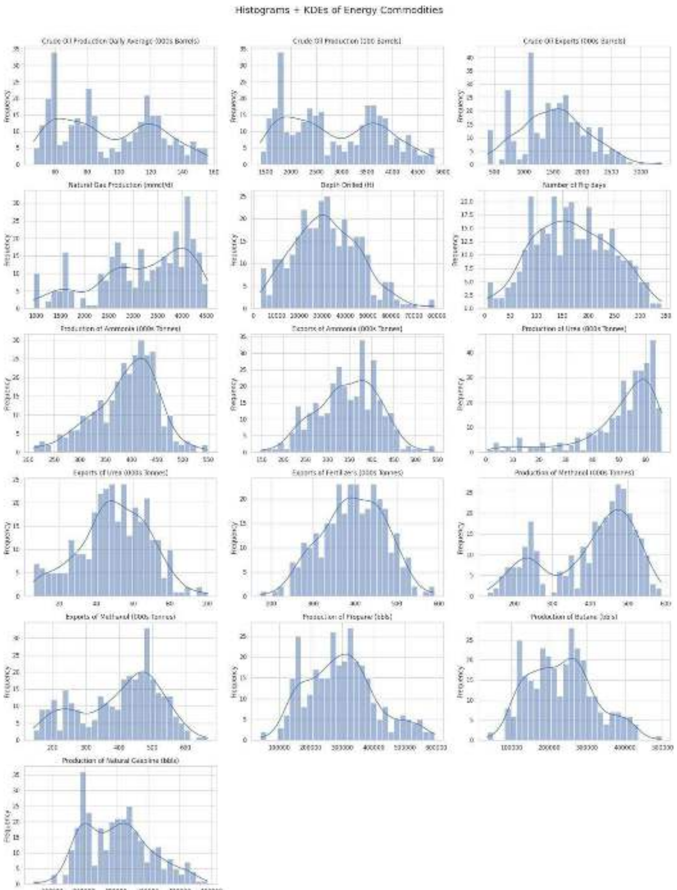

Fig. 1. Distributions of monthly energy commodity indicators with KDE overlays.

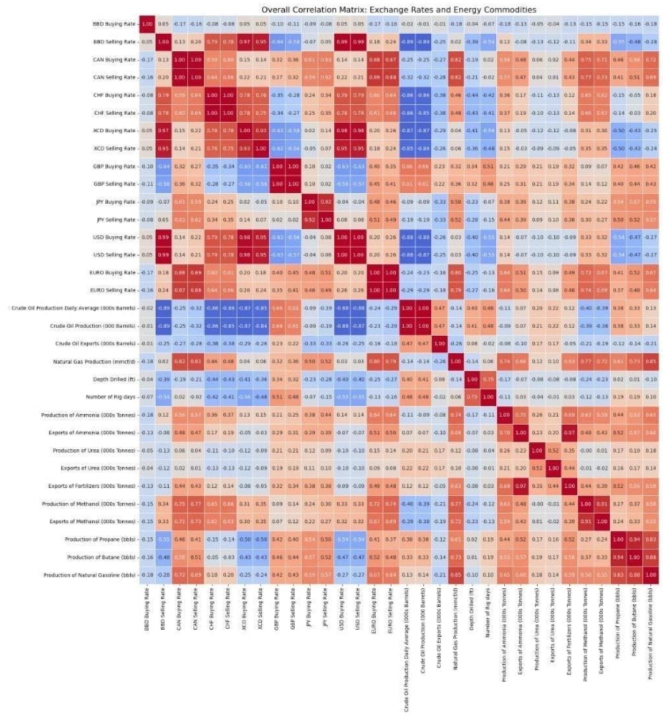

Fig. 2. Correlation heatmap of exchange rates and energy commodity indicators.

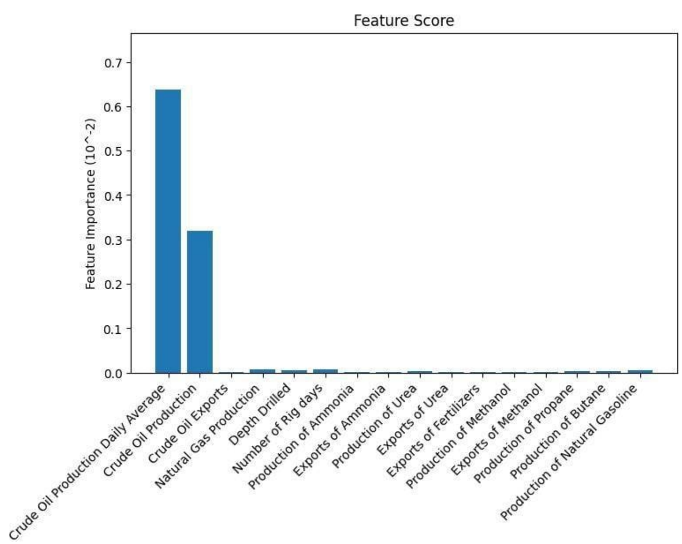

Fig. 3. Feature importance scores for USD buying rate (Random Forest).

not all currencies react similarly to energy sector changes, offering guidance for selecting effective predictors in future models.

B. USD Buying Rate

The USD Buying Rate analysis identified crude oil production daily average (importance score 0.6380) and total crude oil production (importance score 0.3187) as the most influential energy-based predictors (Figure 3). Both reflect the currency's sensitivity to oil market forces. XGBoost testing showed that crude oil production alone produced the lowest MAPE (0.1498%), with both predictors combined yielding a slightly higher error (0.1572%). All predictor-based models significantly outperformed the baseline without predictors (MAPE = 5.5226%), indicating that oil production metrics provide substantial value in short-term USD exchange rate forecasting.

Figure 4 compares model accuracy using MAPE, MAE, and NRMSE. The LSTM+GRU ensemble achieved the lowest MAPE (0.06%) and MAE (0.42%), while LSTM+GRU+XGBoost also performed strongly. This demonstrates the value of combining memory-based and tree-based learning. In contrast, standalone XGBoost and LSTM models yielded weaker results.

These findings highlight the effectiveness of ensemble forecasting for exchange rates influenced by energy variables. By merging the temporal learning strengths of LSTM and GRU, the hybrid models produced precise and stable forecasts, underscoring the potential of integrated approaches for capturing the complex economic drivers of the USD buying rate.

C. EURO Buying Rate

The EURO Buying Rate analysis identified natural gas production (importance score 0.6791) and production of butane

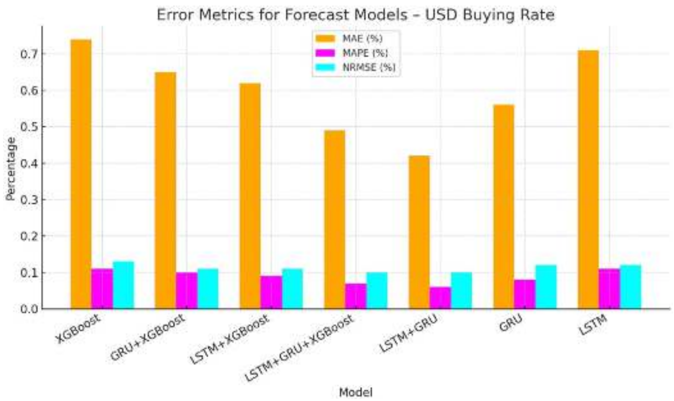

Fig. 4. Error metrics (MAE, MAPE, NRMSE) for forecast models—USD buying rate.

(importance score 0.0724) as the top energy-related predictors (Figure 5). While these factors highlight the petroleum and petrochemical sectors' influence on currency movements, XGBoost tests showed that the baseline model without predictors (MAPE = 3.7265%) achieved better accuracy than models using these features, indicating they did not improve short-term forecasting performance.

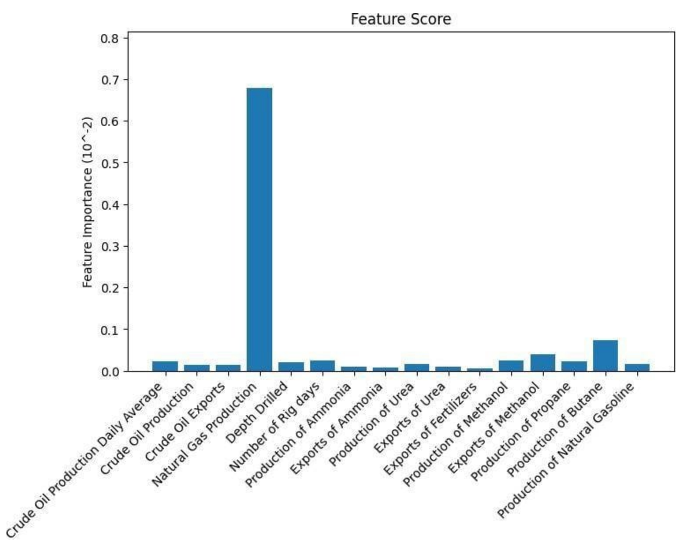

Fig. 5. Feature importance scores for EURO buying rate (Random Forest).

Figure 6 summarizes model performance using MAPE, MAE, and NRMSE. The LSTM+GRU ensemble recorded the lowest MAPE (1.31%) and MAE (9.96%), while GRU and XGBoost had the weakest scores for these metrics. Across all tests, ensemble models outperformed single architectures, with LSTM+GRU consistently ranking best. This suggests that integrating models capable of capturing time dependencies with boosting techniques enhances the tracking of EURO buying rate fluctuations.

D. Summary of Best-Performing Models

Table I summarises the best-performing models for each currency's buying and selling rate, based on the lowest combination of MAPE, MAE, and NRMSE.

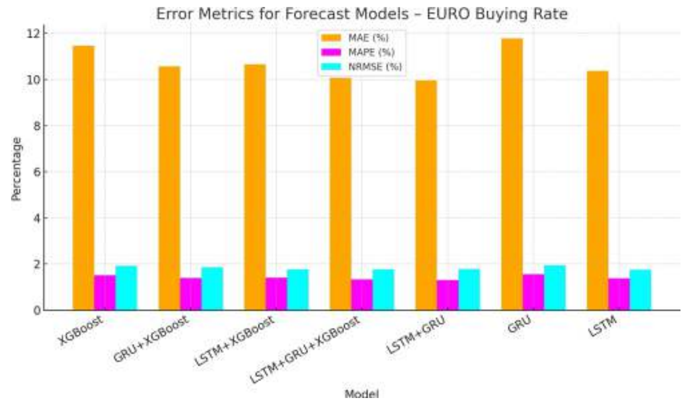

Fig. 6. Error metrics (MAE, MAPE, NRMSE) for forecast models—EURO buying rate.

TABLE I
BEST PERFORMING MODELS FOR EACH CURRENCY RATE

Currency	Rate Type	Best Model
BBD	Buying	LSTM + GRU + XGBoost
BBD	Selling	LSTM + GRU + XGBoost
CAN	Buying	GRU + XGBoost
CAN	Selling	LSTM + XGBoost
CHF	Buying	LSTM + GRU + XGBoost
CHF	Selling	LSTM + GRU + XGBoost
XCD	Buying	LSTM + XGBoost
XCD	Selling	GRU + XGBoost
GBP	Buying	LSTM + XGBoost
GBP	Selling	LSTM + XGBoost
JPY	Buying	GRU
JPY	Selling	GRU
USD	Buying	LSTM + GRU
USD	Selling	XGBoost
EURO	Buying	LSTM + GRU
EURO	Selling	LSTM

V. DISCUSSION

This study evaluated multiple machine learning models for forecasting the monthly buying and selling rates of eight major currencies in 2024 using a rolling origin forecast strategy. The results revealed that ensemble approaches combining temporal sequence learning and nonlinear decision trees, particularly the LSTM+GRU+XGBoost model, frequently outperformed standalone architectures. These hybrid models effectively captured both the sequential dependencies inherent in time-series data and the nonlinear relationships driven by energy market fluctuations. Notably, the GRU model proved most effective for the JPY, underscoring that model suitability can vary by currency due to differing market sensitivities.

Feature selection highlighted crude oil production, methanol exports, and natural gas production as key predictors across currencies. Crude oil, as the cornerstone of Trinidad and Tobago's export economy, influenced several currencies through its effect on foreign currency supply and trade balances. Methanol exports showed targeted impacts, particularly for Barbados and Canada, reflecting specific trade agreements and petrochemical market dynamics. Natural gas production, while not always directly linked to trading partners, contributed to exchange rate movements via its role in foreign exchange reserves and energy revenue streams.

The practical implications of these findings are substantial. For monetary policy, central banks can incorporate commodity-based forecasts into intervention timing and reserve management strategies. Budgetary accuracy can be improved by fiscal planners by including production volumes together with price forecasts to better manage stabilization funds and sovereign wealth. In the energy sector, producers can adjust export schedules, refine hedging strategies, and plan infrastructure investments based on expected currency movements. Additionally, financial institutions may use these projections to plan for liquidity, estimate capital adequacy, and price foreign exchange products, while regional insurers and SMEs that operate cross-border business can make use of lighter, customized versions of the forecasting structure.

VI. Recommendations

While the models succeeded in explaining the relationship between energy production and exchange rate movement, a series of refinements can advance future frameworks. Incorporating techniques for structural break detection such as Zivot–Andrews or Chow tests would consider shifts that are induced by major events such as the 2008 Financial Crisis or COVID-19. Furthermore, the addition of macroeconomic factors such as inflation, interest rates, and fiscal deficit to the model might improve accuracy, particularly for currencies that are less affected by energy. N-BEATS and Transformer models are some of the advanced neural architectures that may be experimented with to better learn short- and long-term relationships. Finally, applying interpretability techniques, including SHAP values or attention mechanisms, would make ensemble deep learning models more transparent for policymakers and financial institutions.

VII. Conclusion

This study developed and evaluated a machine learning–driven framework for forecasting monthly buying and selling rates of eight major currencies against the Trinidad and Tobago dollar. By integrating exchange rate data with energy production statistics, the models examined the influence of commodities such as crude oil, methanol, and natural gas on currency dynamics.

Feature selection revealed crude oil production as the most influential predictor for several currencies, reflecting its pivotal role in Trinidad and Tobago's export economy. Ensemble approaches, particularly LSTM+GRU+XGBoost, consistently delivered superior accuracy over standalone models, confirming the benefit of combining temporal sequence learning with nonlinear decision boundaries for robust short-term forecasting.

The results obtained from this study have numerous applications for monetary policy, fiscal planning, energy sector decision-making, and financial stability management. Central banks can use these forecasts to optimize intervention timing; ministries of finance can enhance revenue planning; energy producers can adjust export and hedging strategies; and financial institutions can incorporate predictions into risk and liquidity management frameworks.

In general, the incorporation of energy indicators into advanced machine learning frameworks offers a valuable decision-support tool for energy-dependent economies. In the context of increased market uncertainty and economic reformation, forward-looking systems like these can increase resilience, promote forward-looking governance, and yield a competitive edge in national and regional financial planning.

References

[1] R. A. Meese and K. Rogoff, "Empirical exchange rate models of the seventies: Do they fit out of sample?" *Journal of International Economics*, vol. 14, no. 1–2, pp. 3–24, 1983.

[2] R. A. Amano and S. van Norden, "Exchange Rates and Oil Prices," *Review of International Economics*, vol. 6, no. 4, pp. 683–694, 1998.

[3] Y.-C. Chen and K. Rogoff, "Commodity currencies," *Journal of International Economics*, vol. 60, no. 1, pp. 133–160, 2003.

[4] G. E. P. Box and G. M. Jenkins, *Time Series Analysis: Forecasting and Control*, Holden-Day, 1970.

[5] S. Makridakis and M. Hibon, "The M3-Competition: Results, conclusions and implications," *International Journal of Forecasting*, vol. 16, no. 4, pp. 451–476, 2000.

[6] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[7] V. Chang, Q. Xu, C. Anyamele, and H. Wang, "Predicting Economic Trends and Stock Market Prices with Deep Learning and Advanced Machine Learning Techniques," *Electronics*, vol. 13, p. 3396, 2024.

[8] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proc. of the 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.

[9] U. Ullah et al., "Forecasting Foreign Exchange Rate with Machine Learning Techniques Chinese Yuan to US Dollar Using XGBoost and LSTM Model," *iRASD Journal of Economics*, vol. 6, pp. 761–775, 2024.

[10] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[11] S. Siami-Namini, N. Tavakoli, and A. Siami Namin, "A Comparison of ARIMA and LSTM in Forecasting Time Series," in *2018 IEEE Int. Conf. on Machine Learning and Applications*, pp. 1394–1401.

[12] K. Cho et al., "Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation," *arXiv preprint arXiv:1406.1078*, 2014.

[13] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "Statistical and Machine Learning forecasting methods: Concerns and ways forward," *PLOS ONE*, vol. 13, no. 3, pp. e0194889, 2018.

[14] R. T. Clemen, "Combining forecasts: A review and annotated bibliography," *International Journal of Forecasting*, vol. 5, no. 4, pp. 559–583, 1989.