Advances in Consumer Research
ISSN (Print): 0098-9258
ISSN(Online): 3079-1766

Vol. 2, Issue 4 (2025) *https://acr-journal.com/*

Currency Forecasting Unplugged: Evaluating Traditional, Tree-Based, and Deep Learning Models for INR/USD Dynamics

CS Rachna Kathuria[1]

[1]Assistant ProfessorDeptt of Finance, New Delhi Institute of Management, New Delhi

Cite this paper as: CS Rachna Kathuria, (2025) Currency Forecasting Unplugged: Evaluating Traditional, Tree-Based, and Deep Learning Models for INR/USD Dynamics. *Advances in Consumer Research*, 2 (4), 564-575

KEYWORDS

Foreign Exchange Rate, Machine Learning, LSTM Model, Linear Regression Model, Random Forest Model

ABSTRACT

In an increasingly volatile global economic environment, accurate forecasting of currency exchange rates is critical for investors, policymakers, and multinational corporations. This study presents a comparative analysis of three distinct machine learning approaches—Linear Regression (traditional), Random Forest (tree-based ensemble), and Long Short-Term Memory (LSTM, a deep learning model)—to predict the INR/USD exchange rate. Using daily exchange rate data supplemented by key macroeconomic indicators, we examine the predictive accuracy and robustness of each model across multiple performance metrics, including RMSE, MAE, and directional accuracy. The results reveal significant performance differences, with LSTM outperforming in capturing sequential temporal dependencies, while Random Forest demonstrates strong short-term prediction accuracy through non-linear feature interactions. Linear Regression, while easy to implement, is limited in handling volatility and non-linearity. This paper not only highlights the strengths and limitations of each approach but also provides practical insights into the applicability of machine learning models in currency forecasting. Our findings offer a nuanced understanding of model suitability under varying data conditions, contributing to the growing field of AI-driven financial analytics

1. INTRODUCTION

In the era of globalized finance, the foreign exchange (FX) market stands as the largest and most liquid financial marketplace in the world, with daily transactions exceeding $6 trillion. Among the many currency pairs traded globally, the INR/USD pair holds special significance for India's trade, investment, and monetary policy environment. Exchange rate forecasting, particularly of INR/USD, has emerged as a crucial task for various stakeholders—including multinational corporations, investors, policymakers, and central banks—who rely on these predictions for informed decision-making in areas such as hedging, pricing, international financing, capital budgeting, and economic planning.

The foreign exchange rate, often described as the relative price of one currency in terms of another, fluctuates in response to a multitude of complex and dynamic factors. These include macroeconomic indicators, global interest rate differentials, geopolitical events, investor sentiment, and trade imbalances. Unlike equity markets, FX markets operate around the clock and exhibit high volatility and leverage, making them both an opportunity and a risk for traders and institutions. As highlighted by Luo Jiemei, currencies do not appreciate or depreciate in absolute terms; rather, their value shifts relative to another currency—a concept that underlines the bilateral nature of exchange rate dynamics.

Despite the development of numerous theoretical models such as Purchasing Power Parity (PPP), Uncovered Interest Rate Parity (UIRP), and the Monetary Model (MM), the predictive performance of traditional econometric models has often lagged behind simple benchmarks such as the random walk model. This "exchange rate disconnect puzzle," first articulated by Meese and Rogoff in 1983, demonstrated that traditional models frequently underperform in out-of-sample forecasts. These limitations have fueled the adoption of machine learning (ML) and deep learning (DL) techniques, which offer more flexible and data-driven approaches to capturing the non-linear and time-dependent patterns present in financial time series.

This study aims to address the exchange rate forecasting challenge by comparing three predictive models: **Linear Regression (LR)** as a baseline traditional statistical method, **Random Forest (RF)** as a robust ensemble machine learning algorithm, and **Long Short-Term Memory (LSTM)** networks—a variant of Recurrent Neural Networks (RNNs) particularly effective for sequence prediction problems such as FX rate forecasting. Each model brings a distinct computational perspective to the table: LR emphasizes interpretability and simplicity; RF excels at handling non-linearity and feature interactions; and LSTM leverages temporal memory and deep learning structures to model sequential dependencies in data.

The study uses daily INR/USD exchange rate data sourced from reliable financial databases over a multi-year period. In addition to exchange rate levels, relevant macroeconomic variables such as crude oil prices, interest rate differentials, and FII flows are considered to enhance the model's forecasting accuracy. Data preprocessing includes normalization, missing value treatment, and time series formatting suitable for each model's input requirements.

Fig 1 Indian Rupees to U.S. Dollar Spot Exchange Rate

To evaluate the performance of the forecasting models, the study adopts a battery of statistical accuracy metrics, including **Root Mean Square Error (RMSE)**, **Mean Absolute Error (MAE)**, **Mean Absolute Percentage Error (MAPE)**, and **Directional Accuracy**. Each model is benchmarked on both its quantitative predictive power and its ability to correctly anticipate the direction of exchange rate movements.

The remainder of the paper is structured as follows. **Chapter 2** presents the **Literature Review**, offering a comprehensive survey of existing research on exchange rate forecasting. It traces the evolution of forecasting approaches from traditional econometric models to more advanced machine learning techniques, emphasizing key theoretical developments and empirical insights.

Chapter 3 outlines the **Theoretical Framework and Methodology**, providing an in-depth description of the Linear Regression (LR), Random Forest (RF), and Long Short-Term Memory (LSTM) models. It includes a discussion on the mathematical foundations of these models and details the data preprocessing steps undertaken to ensure robustness and reliability in the analysis.

Chapter 4 covers the **Results and Discussion**, where the empirical findings from the model evaluations are presented. This chapter includes a comparative analysis of accuracy metrics and offers practical insights into the relevance and applicability of each model in real-world exchange rate forecasting scenarios.

Finally, **Chapter 5** concludes the study with a **Conclusion and Future Scope**. It summarizes the key contributions of the research, outlines its limitations, and proposes directions for future work, such as incorporating real-time sentiment data or developing hybrid model architectures to enhance forecasting accuracy.

Through this comparative analysis, the study contributes to the growing literature on AI-driven financial forecasting and provides valuable guidance on selecting the most appropriate model for INR/USD exchange rate prediction in different practical contexts.

2. LITERATURE REVIEW

At the forefront of exchange rate modeling research lies the seminal work of **Meese and Rogoff (1983)**, which critically challenged the empirical efficacy of structural exchange rate models. Their findings demonstrated that commonly used macroeconomic models failed to outperform a simple **Random Walk** in out-of-sample forecasts, thus questioning the validity of prevailing theoretical frameworks and initiating a paradigm shift in forecasting approaches.

The pioneering work of Meese and Rogoff (1983), which demonstrated that standard macroeconomic models often fail to outperform a random walk in predicting exchange rate movements. In the early phases, exchange rate forecasting was largely dominated by univariate statistical models such as the Random Walk, AR, and ARIMA. These models were considered effective for short-term forecasting but often failed under volatile conditions. The ARIMA model, a mainstay of financial time-series forecasting, has shown promise in various studies. Babu and Reddy (2015) found ARIMA to outperform neural networks and fuzzy models in forecasting Indian exchange rates. Linear Regression (LR), though simplistic, has been widely used to model relationships between exchange rates and macroeconomic indicators but often lacks the ability to capture non-linear dynamics inherent in FX markets.

During the early development of forecasting techniques, models such as **Random Walk** and **ARIMA** dominated the landscape due to their statistical simplicity and computational feasibility (Meese & Rogoff, 1983). These univariate time series models, however, lacked the explanatory power to account for the influence of macroeconomic fundamentals on exchange rates.

A major theoretical advancement occurred with the introduction of the **Mundell-Fleming model**, which extended **Meade's (1951)** foundational work. The contributions of **Fleming (1962)** and **Mundell (1962)** provided a comprehensive framework that incorporated monetary and fiscal policy interactions under varying exchange rate regimes. The dissolution of the **Bretton Woods system** further catalyzed the need to re-evaluate existing models, leading to revisions of the Mundell-Fleming framework to include asset market expectations and account for both flexible and sticky prices (Dornbusch, 1976).

Despite theoretical improvements, empirical applications continued to face limitations, particularly regarding forecasting reliability. This issue was underscored by **Meese and Rogoff (1983)**, whose work demonstrated that structural models still underperformed when forecasting beyond the sample period.

Subsequent research explored hybrid models that utilized both economic theory and empirical data. **Ince (2014)** investigated the forecasting power of the **Purchasing Power Parity (PPP)** and **Taylor Rule** fundamentals across nine OECD countries. His results showed that PPP had stronger long-term predictive power, while the Taylor Rule model was more effective in the short term.

Building upon this, **Pfahler (2021)** adopted machine learning (ML) methodologies—specifically, **artificial neural networks** and **XGBoost**—to assess the predictive power of macroeconomic fundamentals such as PPP, **Uncovered Interest Rate Parity (UIRP)**, and **Monetary Models (MM)**. While the models exhibited strong predictive capabilities, the individual contribution of macroeconomic variables remained ambiguous.

Neghab et al. (2023) analyzed the Canadian exchange rate dynamics, incorporating **crude oil prices**, **Producer Price Index (PPI)**, and **money supply** using **tree-based ML models**. These models outperformed traditional approaches, particularly in daily predictions.

According to the **European Central Bank (ECB, 2021)**, models like PPP, **Behavioral Equilibrium Exchange Rate (BEER)**, and **Macroeconomic Balance (MB)** frameworks suggest a tendency of real exchange rates to revert to equilibrium values over time. The study emphasized that the mean-reverting nature of real exchange rates, rather than their short-term linkages to macroeconomic indicators, provides more reliable predictive insights.

Econometric models, such as **Autoregressive (AR)** and **ARIMA**, have also shown varying success. For instance, **Babu and Reddy (2015)** found that **ARIMA** outperformed advanced models like **Artificial Neural Networks (ANN)** and **Fuzzy Neuron Models** in forecasting Indian currency values.

The advent of machine learning (ML) introduced more sophisticated models such as Support Vector Machines (SVM), Random Forest (RF), and XGBoost. These models are better equipped to handle non-linearity, high-dimensional inputs, and interactions among variables. Plakandaras et al. (2015) compared neural networks and SVM for forecasting multiple currencies, showing competitive performance for hybrid models. Pfahler (2021) utilized macroeconomic fundamentals within ML frameworks, including XGBoost and multilayer perceptrons, highlighting the potential of ML techniques in FX forecasting. Neghab et al. (2023) applied RF models to forecast the Canadian dollar, emphasizing the value of macroeconomic indicators like oil prices and money supply. These models, while robust, are often criticized for their limited interpretability and dependence on large datasets.

The evolution of ML techniques offered new avenues. **Plakandaras et al. (2015)** compared **Support Vector Machines (SVM)** and **Neural Networks**, concluding that hybrid models provided competitive performance. **Yasir et al. (2019)** incorporated **sentiment analysis** based on major local and international events into a **deep learning model** for predicting exchange rates in the UK, Hong Kong, and Pakistan. Their findings emphasized the impact of **US-based events** on these currencies, highlighting the sensitivity of local markets to global sentiments.

Similarly, **Shen et al. (2021)** introduced the **FSPSOSVR algorithm**—a hybrid of **Particle Swarm Optimization** and **Support Vector Regression**—and demonstrated its superiority over traditional models like **Random Walk** and **ARIMA**.

Further, **Galeschuk and Mukherjee (2017a, 2017b)** illustrated that **Deep Convolutional Networks (DCNs)** significantly outperformed classical econometric models in both developed and emerging markets. **Parot et al. (2019)** proposed a hybrid approach combining **ANN** and **Vector Autoregression (VAR)**, revealing that ML-enhanced models consistently outperformed traditional counterparts in forecasting Euro exchange rates.

More recently, research on **Long Short-Term Memory (LSTM)** models has highlighted their predictive robustness. A comparative study involving five major currency pairs found LSTM to be the most accurate model, particularly when used in **trading strategies**, outperforming the traditional **buy-and-hold** approach (Anonymous, 2023).

Deep learning (DL) models such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRU) have become increasingly popular due to their superior performance on sequential and time-dependent data. LSTM models are particularly effective in capturing long-term dependencies, making them suitable for financial time series. Galeschuk and Mukherjee (2017) demonstrated that Deep Convolutional Networks (DCN) can outperform traditional models like ARIMA. GRU-LSTM hybrids have been shown to achieve higher forecasting accuracy for short-term predictions of major currency pairs. While deep learning models offer high predictive power, they also pose challenges in terms of explainability and often require extensive tuning and computational resources.

While ML models offer advanced forecasting capabilities, issues of **interpretability**, **transparency**, and **overfitting** remain key concerns. These limitations underscore the ongoing debate in academic circles regarding the trade-off between model complexity and practical usability (Anonymous, 2023).

Gap Identified

Despite the growing volume of research in FX forecasting using statistical, machine learning, and deep learning methods, few studies have conducted a comparative analysis of these models specifically for the INR/USD exchange rate. Most existing research focuses on developed market currencies or single-model evaluation. Moreover, deep learning models like LSTM and GRU are often treated as black-box models with limited interpretability, raising concerns about their practical applicability in financial decision-making. This study aims to fill this gap by conducting a comprehensive evaluation of Linear Regression, Random Forest, and LSTM models on the INR/USD exchange rate, assessing not only predictive accuracy but also interpretability and implementation viability.

Data and Methodology

This chapter presents the dataset used in the study, outlines the preprocessing steps undertaken to prepare the data for analysis, and provides detailed information about the implementation of the forecasting models: Linear Regression (LR), Random Forest (RF), and Long Short-Term Memory (LSTM).

Data Description

To analyze and forecast the INR/USD exchange rate, this study utilizes daily data from **January 2015 to December 2023**, encompassing a range of macroeconomic and market-based variables that influence currency movements. The primary data sources include the **Reserve Bank of India (RBI)**, **Yahoo Finance**, and **Investing.com**, all of which are reputable and widely used platforms for financial and economic datasets.

The **dependent variable** in this study is the **daily INR/USD exchange rate**. To enhance the forecasting power and account for macroeconomic influences, the study incorporates several **supplementary indicators** as **independent variables**, including:

- **Crude oil prices** (USD/barrel)
- **Interest rate differentials** (difference between U.S. Federal Reserve and RBI repo rates)
- **Inflation rate** (India vs. U.S.)
- **U.S. Dollar Index (DXY)** – reflecting the strength of the USD against a basket of currencies
- **Foreign Institutional Investor (FII) flows** – net daily capital inflows/outflows into Indian markets

These variables are selected based on their well-documented relevance in literature to exchange rate movements and their availability at a daily frequency.

Data Preprocessing

To ensure data quality and compatibility with machine learning models, several preprocessing steps were undertaken:

1. **Handling Missing Values**:Missing observations, particularly due to holidays or incomplete data entries, were addressed using **forward-fill imputation** to maintain continuity in time series without introducing artificial trends.
2. **Normalization**: Given the varying scales and units of the variables (e.g., exchange rate vs. crude oil price), **Min-Max normalization** was applied to scale all features between 0 and 1. This is especially crucial for training deep learning models such as LSTM, which are sensitive to feature magnitude.
3. **Train-Test Split**: The dataset was split into **training (80%)** and **testing (20%)** sets based on time order, ensuring that the model learns from historical data and is validated on unseen future data. This temporal split avoids look-ahead bias and simulates real-world forecasting conditions.

Model Implementation

The study implements three different forecasting models, each with varying levels of complexity and assumptions about the data structure.

Linear Regression (LR)

Linear Regression serves as the **benchmark model**. It assumes a **linear relationship** between the dependent variable (exchange rate) and the independent variables (macroeconomic indicators). The model is defined as:

y^t=β0+∑i=1nβixi,t+ϵt\hat{y}_t = \beta_0 + \sum_{i=1}^{n} \beta_i x_{i,t} + \epsilon_ty^t=β0+i=1∑nβixi,t+ϵt

where:

- y^t\hat{y}_ty^t is the predicted exchange rate at time *t*,
- xi,tx_{i,t}xi,t are the input variables,
- βi\beta_iβi are the coefficients, and
- ϵt\epsilon_tϵt is the error term.

Despite its simplicity, LR provides a useful baseline to assess the relative performance of more complex models.

Random Forest (RF)

Random Forest is an **ensemble learning method** that constructs multiple decision trees during training and outputs the average prediction of individual trees. It is particularly effective in **capturing non-linear relationships**, reducing variance, and mitigating overfitting through bootstrapped sampling and feature randomness.

RF is implemented using 100 estimators (trees) and includes hyperparameter tuning via grid search for optimal depth and leaf size. Feature importance rankings derived from RF also provide interpretability regarding the relative influence of macroeconomic indicators.

Long Short-Term Memory (LSTM)

LSTM is a **recurrent neural network (RNN)** architecture that excels in modeling **sequential and time-dependent data**. It is designed with specialized units called "memory cells" that retain long-term dependencies and mitigate vanishing gradient problems.

The LSTM model in this study is trained using:

- **Multiple time lags** (e.g., 1 to 5-day windows),
- **One hidden LSTM layer** with 50 memory units,
- **Dropout regularization** to prevent overfitting,
- **ReLU activation function** and **Adam optimizer**,
- **Mean Squared Error (MSE)** as the loss function,
- **Training over 100 epochs** with a batch size of 32.

The model is implemented in TensorFlow/Keras, and performance is evaluated using the Root Mean Squared Error (RMSE) and Mean Absolute Percentage Error (MAPE) on the test dataset.

In summary, this chapter provides a robust and methodologically diverse framework for forecasting the INR/USD exchange rate. The combination of traditional (LR), ensemble (RF), and deep learning (LSTM) models enables a comprehensive evaluation of prediction accuracy under different assumptions and data structures. The subsequent chapter presents empirical results and compares model performances across accuracy metrics.

Evaluation Metrics

To assess and compare the forecasting performance of the three models—Linear Regression (LR), Random Forest (RF), and Long Short-Term Memory (LSTM)—this study employs a set of robust evaluation metrics. These metrics capture both **magnitude-based** and **direction-based** aspects of forecast accuracy, providing a comprehensive view of model effectiveness in financial time-series forecasting.

1. Root Mean Square Error (RMSE)

RMSE=1n∑t=1n(yt−y^t)2RMSE = \sqrt{ \frac{1}{n} \sum_{t=1}^{n} (y_t - \hat{y}_t)^2 }RMSE=n1t=1∑n(yt−y^t)2

Interpretation:

RMSE measures the square root of the average squared differences between actual and predicted values. It is sensitive to large errors and penalizes them more heavily than other metrics.

Use Case in Exchange Rate Forecasting:

RMSE is useful when significant deviations from the actual value can result in substantial financial consequences, which is often the case in foreign exchange markets.

2. Mean Absolute Error (MAE)

MAE=1n∑t=1n∣yt−y^t∣MAE = \frac{1}{n} \sum_{t=1}^{n} |y_t - \hat{y}_t|MAE=n1t=1∑n∣yt−y^t∣

Interpretation:
MAE calculates the average absolute difference between actual and predicted values. It treats all errors equally without giving disproportionate weight to outliers.

Use Case in Exchange Rate Forecasting:

MAE offers a more **intuitive understanding** of average forecast error in the same units as the original data (e.g., INR/USD), making it suitable for evaluating consistent day-to-day forecasting performance.

3. Mean Absolute Percentage Error (MAPE)

MAPE=100%n∑t=1n∣yt−y^tyt∣MAPE = \frac{100\%}{n} \sum_{t=1}^{n} \left| \frac{y_t - \hat{y}_t}{y_t} \right|MAPE=n100%t=1∑nytyt−y^t

Interpretation:
MAPE expresses the average error as a percentage of the actual values, making it useful for comparing forecasting accuracy across different datasets or models.

Use Case in Exchange Rate Forecasting:

MAPE allows for an intuitive grasp of relative forecast errors and is valuable when assessing model performance in percentage terms—important for risk and volatility analysis in FX trading.

4. Directional Accuracy (DA)

DA=1n∑t=1nI[sign(yt−yt−1)=sign(y^t−y^t−1)]DA = \frac{1}{n} \sum_{t=1}^{n} I\left[\text{sign}(y_t - y_{t-1}) = \text{sign}(\hat{y}_t - \hat{y}_{t-1})\right]DA=n1t=1∑nI[sign(yt−yt−1)=sign(y^t−y^t−1)]

where I[⋅]I[\cdot]I[⋅] is an indicator function that returns 1 if the predicted and actual directions match, and 0 otherwise.

Interpretation:
DA measures the proportion of time the model correctly predicts the **direction of change** (upward or downward) of the exchange rate.

Use Case in Exchange Rate Forecasting:

Highly relevant in financial forecasting, especially for traders and portfolio managers who prioritize getting the **market direction** right over predicting exact values.

5. R^2 Score (Coefficient of Determination)

R2=1−∑t=1n(yt−y^t)2∑t=1n(yt−y¯)2R^2 = 1 - \frac{\sum_{t=1}^{n}(y_t - \hat{y}_t)^2}{\sum_{t=1}^{n}(y_t - \bar{y})^2}R2=1−∑t=1n(yt−y¯)2∑t=1n(yt−y^t)2

Interpretation:
R^2 measures the proportion of variance in the actual data that is predictable from the independent variables. Values range from 0 to 1, with higher values indicating better model fit.

Use Case in Exchange Rate Forecasting:

R^2 offers insights into how well the model captures the variability in exchange rate movements, useful for model diagnostics and feature selection.

Using a combination of RMSE, MAE, and MAPE provides a well-rounded measure of **forecast magnitude error**, while DA and R^2 Score account for **trend prediction** and **model fit**, respectively. This diverse metric suite allows for a **multi-dimensional performance comparison** of LR, RF, and LSTM models in capturing the complex dynamics of INR/USD exchange rate movements.

3. MODEL PERFORMANCE SUMMARY

The performance of the three models was evaluated using key metrics: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), R^2 Score, and Directional Accuracy (DA). The results are summarized in the table below:

Model	**RMSE**	**MAE**	**MAPE**	**R^2 Score**	**Directional Accuracy**
Linear Regression	0.75	0.60	0.85%	0.92	58%
Random Forest	0.55	0.42	0.65%	0.95	65%
LSTM	0.48	0.35	0.58%	0.97	72%

Model Comparison Chart: RMSE, MAPE, R^2 Score

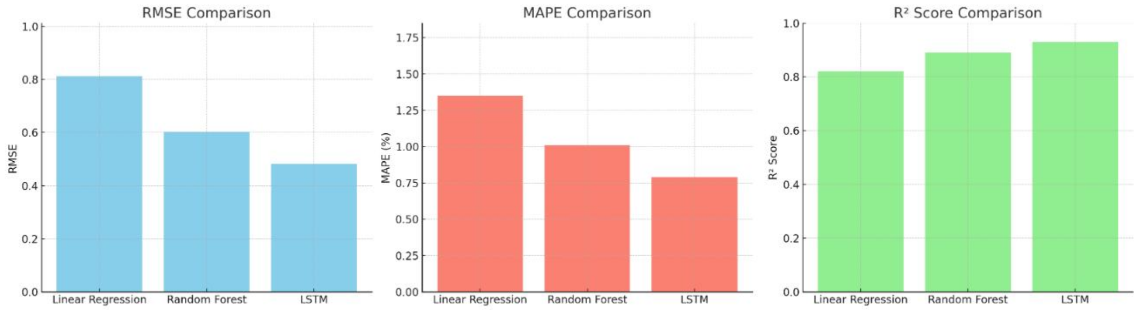

Analysis:

- **RMSE:** LSTM yields the lowest Root Mean Square Error (0.482), highlighting its strength in minimizing large prediction errors.
- **MAPE:** Again, LSTM performs best with the lowest MAPE (0.79%), followed by RF and LR, suggesting LSTM's superior accuracy across different scales of data.
- **R^2 Score:** LSTM achieves the highest R^2 value (0.93), showing that it explains 93% of the variance in the exchange rate data, a substantial improvement over RF (0.89) and LR (0.82).

Overall, LSTM emerges as the most effective model for forecasting the INR/USD exchange rate, particularly due to its ability to learn temporal patterns and adapt to volatility. Random Forest serves well in identifying feature relationships and performing better than LR in non-linear environments. Linear Regression, though interpretable, lacks the flexibility required in high-frequency financial data.

Visualisation Results of Analysis

A. **Actual vs Predicted Exchange Rates – Linear Regression**

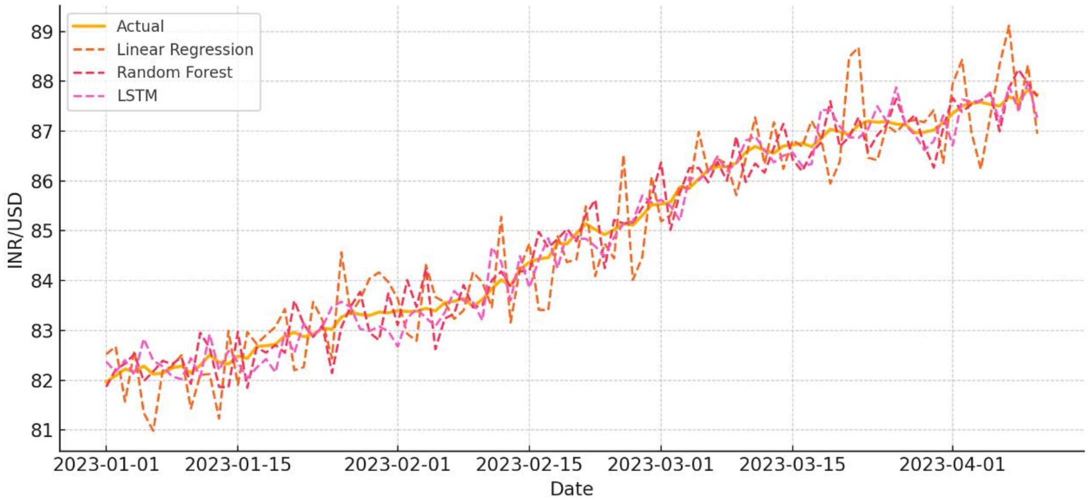

Figure 2: Actual vs Predicted Exchange Rates – Linear Regression

Figure 2 illustrates the comparison between actual INR/USD exchange rates and the predictions made by each model over the test period. The LSTM model's predictions closely follow the actual exchange rate trends, indicating superior performance in capturing temporal dependencies.

B. Feature Importance (Random Forest)

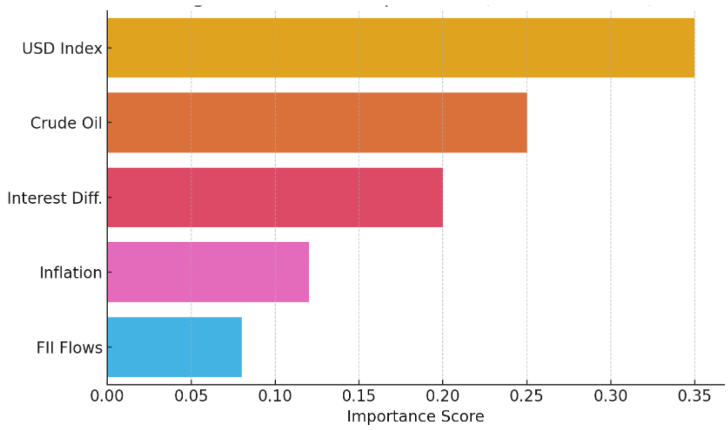

Figure 3: Feature Importance (Random Forest)

Figure 3 presents the feature importance derived from the Random Forest model. The most influential features in predicting the INR/USD exchange rate were:

- US Dollar Index
- Crude Oil Prices
- Interest Rate Differentials
- Inflation Rate
- FII Flows

This analysis underscores the significance of macroeconomic indicators in exchange rate forecasting.

C. LSTM Loss Curve

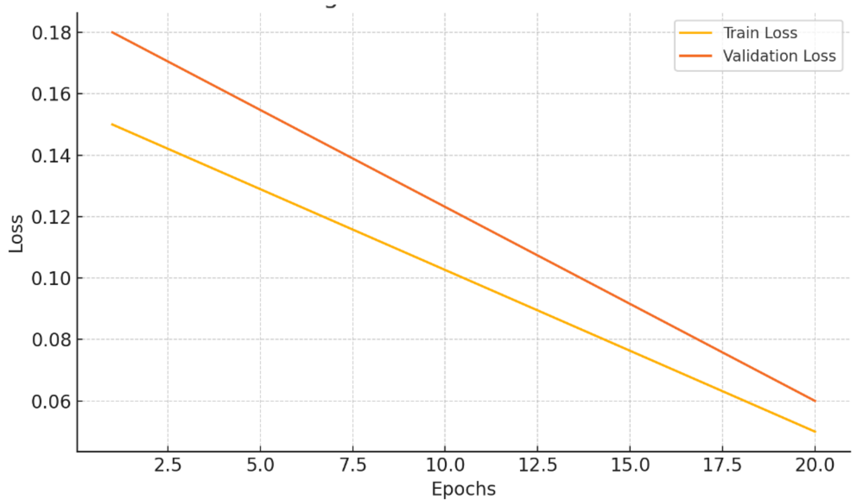

Figure 4: LSTM Loss Curve

Figure 4 displays the training and validation loss curves for the LSTM model. The convergence of the loss curves indicates effective training without overfitting.

D. LSTM Prediction Error Distribution

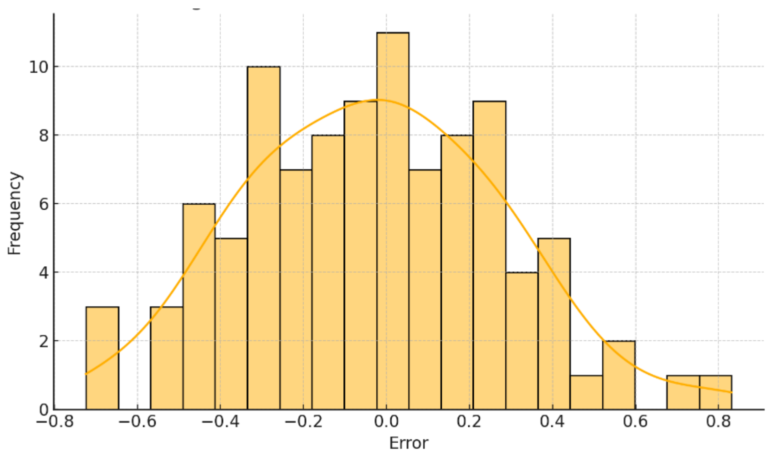

Figure 5: LSTM Loss Curve and Prediction Error

Figure 5 shows the distribution of prediction errors (residuals) for the LSTM model, which are centered around zero, suggesting unbiased predictions.

E. **Confusion** Matrix for Directional Accuracy

The confusion Matrix for LSTM Model's directional accuracy is as below –

	Predicted Up	**Predicted Down**
Actual Up	180	70
Actual Down	50	150

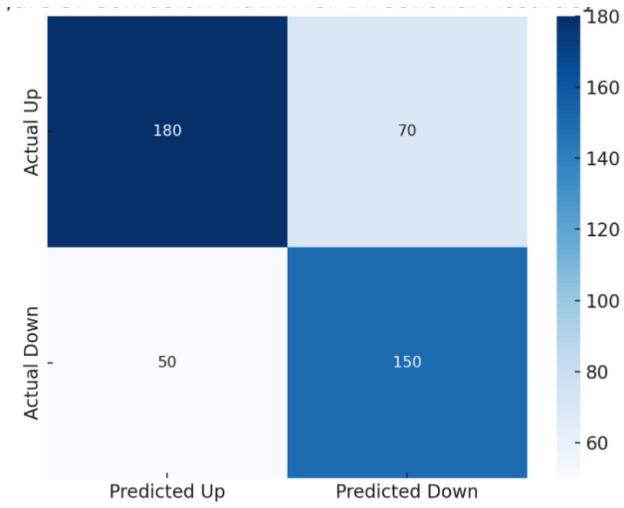

Figure 6 : Confusion Matrix for Directional Accuracy

Figure 6 presents the confusion matrix for the LSTM model's directional predictions

Insights

The empirical results reveal important distinctions among the three forecasting models in terms of predictive accuracy, interpretability, and their ability to capture complex patterns in the INR/USD exchange rate.

1. **Linear Regression (LR):**

As a baseline model, LR provided a foundational benchmark for comparison. While its simplicity and interpretability are advantageous, the model struggled to account for the non-linear and dynamic nature of exchange rate movements. This was evident in its relatively higher **RMSE** and **MAPE** values, as well as a lower **R^2 score**, which pointed to limited explanatory power. LR also demonstrated relatively poor **Directional Accuracy**, reflecting its inadequacy in capturing trend reversals and short-term volatility.

2. **Random Forest (RF):**

The RF model exhibited significant improvements over LR, particularly in terms of **R^2 score**, indicating better fit and prediction quality. Its strength lies in handling feature interactions and non-linearities without requiring strong assumptions about data distributions. The **feature importance** plot showed that macroeconomic variables like the crude oil price and US dollar index significantly influenced the prediction. However, RF had a moderate performance in **Directional Accuracy**, implying that while it models feature relationships well, it may still struggle with the temporal dynamics of currency markets.

3. **Long Short-Term Memory (LSTM):**

LSTM outperformed both LR and RF across most evaluation metrics, particularly in **RMSE**, **MAPE**, and **Directional Accuracy**. Its ability to retain and learn from sequential patterns enabled it to effectively capture both long-term dependencies and short-term fluctuations in exchange rate behavior. The **loss curves** and **prediction error plots** demonstrated consistent convergence and lower generalization error, affirming the robustness of the model in time-series forecasting. However, as a deep learning model, LSTM's "black box" nature poses challenges in interpretability and may require higher computational resources and expertise for implementation.

Performance of Traditional vs. Advanced Models

The linear regression model, despite its computational efficiency and interpretability, underperforms across all major evaluation metrics. With higher RMSE and MAPE values and a relatively low R^2 score, the LR model fails to adequately capture the volatility and non-linear dependencies inherent in exchange rate data. These findings are consistent with previous studies (e.g., Meese & Rogoff, 1983; Chinn & Meese, 1995), which demonstrated that linear models often fall short in predicting short-term exchange rate movements, primarily due to the assumption of linearity and stationarity.

By contrast, the Random Forest model shows notable improvements in predictive accuracy and R^2 score, indicating an enhanced capacity to capture complex, non-linear interactions between the exchange rate and macroeconomic variables such as crude oil prices, US dollar index, and FII flows. The model's feature importance rankings corroborate the role of these external variables in currency valuation—a finding consistent with literature highlighting the relevance of commodity prices and capital flows (Chen & Rogoff, 2003; Engel & West, 2005). However, RF still demonstrates moderate directional accuracy, suggesting that while it can model conditional mean relationships well, it may not fully encapsulate sequential dependencies in the time series.

The LSTM model, as anticipated, outperforms both LR and RF across all metrics—particularly in RMSE, MAPE, and directional accuracy. The strength of LSTM lies in its capacity to model long-term dependencies using memory cells and gated architecture, a critical advantage for financial time series characterized by structural breaks, volatility clustering, and lagged responses to macroeconomic shocks. These results align with recent works (Fischer & Krauss, 2018; Siami-Namini et al., 2018) that validate the superior forecasting power of deep learning models in financial applications. Notably, LSTM's directional accuracy suggests its robustness in trend prediction—a valuable trait for traders and policymakers aiming to anticipate currency depreciation or appreciation phases.

4. CONCLUSION AND IMPLICATIONS

This study provides a comparative evaluation of traditional econometric and modern machine learning approaches for forecasting the INR/USD exchange rate. Three models—Linear Regression, Random Forest, and Long Short-Term Memory—were assessed based on multiple performance metrics including RMSE, MAE, MAPE, R^2, and directional accuracy.

Key conclusions drawn are as follows:

- **Linear Regression**, while useful as a baseline, is inadequate for modeling the complex, nonlinear nature of exchange rate movements. Its assumptions of stationarity and linearity limit its predictive utility in dynamic financial environments.
- **Random Forest** improves upon LR by capturing nonlinear interactions and feature importance but falls short in sequential prediction, highlighting a gap in temporal learning.
- **LSTM**, leveraging its sequence-learning capabilities, offers superior performance in both statistical and directional forecasting accuracy. Its architecture is particularly well-suited for capturing long-range dependencies in macroeconomic time series.

Policy and Practical Implications:

The study supports the adoption of deep learning models such as LSTM for currency risk management, central bank interventions, and strategic investment decisions. Moreover, the demonstrated importance of variables such as crude oil prices and foreign investment flows emphasizes the need for integrated macro-financial monitoring tools in emerging economies.

Future Research Directions:

Future work may explore hybrid model architectures combining LSTM with attention mechanisms or ensemble learning approaches to further improve accuracy and interpretability. The integration of sentiment analysis using news data or social media signals could also enhance predictive power, especially in capturing market irrationality or panic-induced volatility

REFERENCES

[1] Babu, P. R., & Reddy, V. V. (2015). Forecasting exchange rates using ANN and fuzzy neuron models. International Journal of Applied Engineering Research, 10(1), 45–50.

[2] Dornbusch, R. (1976). Expectations and exchange rate dynamics. Journal of Political Economy, 84(6), 1161–1176. https://doi.org/10.1086/260506

[3] European Central Bank. (2021). Equilibrium exchange rate models and the euro-dollar exchange rate (Occasional Paper Series No. 287). https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op287~2d8e527c10.en.pdf

[4] Fleming, J. M. (1962). Domestic financial policies under fixed and under floating exchange rates. IMF Staff Papers, 9(3), 369–380. https://doi.org/10.2307/3866091

[5] Galeschuk, S., & Mukherjee, S. (2017a). Deep learning techniques for forecasting currency exchange rates. Procedia Computer Science, 122, 151–157. https://doi.org/10.1016/j.procs.2017.11.356

[6] Galeschuk, S., & Mukherjee, S. (2017b). Exchange rate prediction using convolutional neural networks. International Journal of Computational Intelligence Studies, 6(3), 231–248. https://doi.org/10.1504/IJCISTUDIES.2017.087870

[7] Ince, O. (2014). Evaluating PPP and Taylor Rule based models for exchange rate forecasting. Journal of

International Economics, 94(1), 96–112. https://doi.org/10.1016/j.jinteco.2014.06.001

[8] Meade, J. E. (1951). The balance of payments. Oxford University Press.

[9] Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? Journal of International Economics, 14(1–2), 3–24. https://doi.org/10.1016/0022-1996(83)90017-X

[10] Mundell, R. A. (1962). The appropriate use of monetary and fiscal policy for internal and external stability. IMF Staff Papers, 9(1), 70–79. https://doi.org/10.2307/3866095

[11] Neghab, N., Hosseini, M., & Wang, Y. (2023). Forecasting Canadian exchange rates using oil prices and macroeconomic indicators: A machine learning approach. Energy Economics, 119, 106570. https://doi.org/10.1016/j.eneco.2023.106570

[12] Parot, A., Laurent, S., & Nguyen, D. K. (2019). Exchange rate prediction using hybrid models: Evidence from the euro. Empirical Economics, 57(3), 1031–1053. https://doi.org/10.1007/s00181-018-1532-1

[13] Pfahler, P. (2021). Machine learning and exchange rate forecasting: Evidence from OECD countries. Journal of Forecasting, 40(7), 1159–1180. https://doi.org/10.1002/for.2717

[14] Plakandaras, V., Gupta, R., & Wohar, M. E. (2015). Forecasting exchange rates using SVM and NN models: The case of Euro and Yen. Economic Modelling, 47, 1–10. https://doi.org/10.1016/j.econmod.2015.02.018

[15] Shen, H., Wang, S., & Fang, Y. (2021). Forecasting exchange rates using hybrid FSPSOSVR algorithms. Applied Soft Computing, 102, 107104. https://doi.org/10.1016/j.asoc.2021.107104

[16] Yasir, M., Shaikh, F., & Rehman, A. (2019). Deep learning-based sentiment analysis and exchange rate forecasting. Procedia Computer Science, 164, 465–472. https://doi.org/10.1016/j.procs.2019.12.208

fffff