Nepal Journal of Mathematical Sciences (NJMS)
ISSN: **2738-9928** (*online*), **2738-9812** (*print*)
Vol. 6, No. 1, 2025 (February): 7-20
DOI: 10.3126/njmathsci.v6i1.77369

Research Article
Received Date: February 10, 2025
Accepted Date: March 20, 2025
Published Date: April 8, 2025

Analysis of Foreign Exchange Rate Forecasting of Nepal using Long Short-Term Memory and Gated Recurrent Unit

Nissan Neupane[1] & Nawaraj Paudel[2*]

[1]*School of Mathematical Sciences, Tribhuvan University, Kirtipur, Kathmandu, Nepal*

[2]*Central Dept. of Computer Science and Information Technology, Tribhuvan University, Kathmandu, Nepal*

Corresponding Author: *nawarajpaudel@cdcsit.edu.np

Abstract: *Foreign exchange rate represents the value of one currency relative to another and influences international trade and investment. It is crucial for a country's economy as it affects the cost of imports and exports, impacting trade balances and inflation rates. This study compares the forecasting of the forex rate of Nepal and its volatility by using Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) models. The study uses secondary time series data that consists of foreign exchange rate from 2005 to 2024 A.D. Various error metrics were used to compare the performance of these models to predict the foreign exchange rate. The final result showed that the LSTM model outperformed GRU with superior forecasting accuracy, achieving a Mean Squared Error (MSE) of 4.7056, a Root Mean Squared Error (RMSE) of 2.1692, a Mean Absolute Error (MAE) of 2.0262, and a Mean Absolute Percentage Error (MAPE) of 1.5764%. In contrast, the GRU model yielded higher error metrics with an MSE of 7.1607, RMSE of 2.6759, MAE of 2.5673, and MAPE of 2.0061%. These findings highlight the effectiveness of LSTM in capturing historical trends and managing volatility, suggesting its robustness for forex rate prediction. Although the study focused on historical forex rates of the Nepalese Rupee against the US Dollar, incorporating additional economic indicators such as interest rates and Foreign Direct Investment (FDI) could enhance the model's predictive capabilities.*

Keywords: Forex rate, Long Short-Term Memory, Gated Recurrent Unit

1. Introduction

In today's globalized economy, the exchange rate between currencies plays a crucial role in shaping the economic landscape of a nation. The foreign exchange market is a global marketplace for exchanging national currencies against one another, and it is the largest and most liquid financial market in the world. Exchange rates fluctuate constantly due to various factors, including economic indicators, geopolitical events, and market speculation. For countries like Nepal, where the economy is significantly influenced by international trade and remittances, the exchange rate between the Nepalese Rupee and the US Dollar is particularly significant. This exchange rate impacts economy in multiple ways. Since Nepal is heavily dependent on imports, particularly from India and other countries where transactions are often denominated in USD, any fluctuations in the exchange rate can affect the cost of goods and services in the country. A depreciation of the NPR against the USD can lead to higher import costs, which may contribute to inflationary pressures within the domestic economy. Conversely, an appreciation of the NPR could make imports cheaper but might reduce the competitiveness of Nepalese exports, thereby affecting the trade balance. These dynamics underscore the importance of accurate forex rate forecasting, which can help policymakers, businesses, and investors make informed decisions [13].

Traditional forex forecasting methods like Autoregressive Integrated Moving Average and Vector Autoregression have been widely used, relying on historical data to identify trends. However, these linear models struggle to capture the non-linear relationships and complex dependencies typical in financial time series, limiting their forecasting accuracy [2].

In response to the limitations of traditional forecasting methods, the use of machine learning and deep learning techniques has gained significant attention in recent years. Machine learning models, particularly Neural Networks have shown great potential in forecasting complex time series data. GRU networks are a type of Recurrent Neural Network which use gating mechanisms to control the flow of information, helping to overcome the issues of exploding gradients commonly encountered in traditional RNNs. GRUs are particularly valuable for tasks such as time series forecasting and financial predictions, where capturing historical dependencies can significantly improve accuracy [4]. LSTM networks, another type of Recurrent Neural Network, are designed to address the problem of long-term dependencies in time series data. Traditional RNNs struggle with the vanishing gradient problem, where the gradients used to update the weights of the network become extremely small, leading to difficulties in learning long-term dependencies. LSTMs overcome this issue by introducing memory cells and gates that control the flow of information, allowing the network to retain important information over long sequences and thus providing more accurate forecasts [9]. The application of LSTM and GRU models to forex rate forecasting has been explored in various studies with promising results. Authors in [1] suggested that LSTM and Artificial Neural Network are the most commonly used machine learning algorithms for forex market. Similarly, authors in [7] highlighted the effectiveness of LSTM networks in financial forecasting, noting their ability to model the sequential nature of time series data more effectively than traditional models. Despite extensive research on machine learning, limited work has focused on emerging markets like Nepal. Most studies target developed economies with different data and market conditions. This study addresses that gap by comparing LSTM and GRU models in forecasting Nepal's USD/NPR exchange rate.

1.1. Research Questions

The specific research question that is addressed is as follows:

- How do LSTM and GRU models forecast foreign exchange rate of Nepalese Rupee against US Dollar?
- What is the comparative performance of LSTM and GRU models in predicting the Nepalese Rupee against US Dollar?

1.2. Research Objectives

The objectives of this study are:

- To develop and implement LSTM and GRU models for forecasting Nepalese Rupee against US Dollar.
- To evaluate and compare the prediction performance of LSTM and GRU models in forecasting Nepalese Rupee against US Dollar.

2. Literature Review

Over the years, various models and approaches have been developed to predict exchange rate movements, ranging from traditional economic theories such as purchasing power parity (PPP) and interest rate parity (IRP) to more advanced econometric and machine learning techniques. Early studies, such as those by [10], challenged the predictive power of fundamental models, emphasizing the superiority of random walk behavior. Since then, researchers have explored time series models, artificial intelligence, and hybrid approaches to improve forecasting accuracy. Despite these advancements, exchange rate prediction remains a complex and debated field due to the influence of multiple economic, political, and speculative factors. Numerous experiments have been carried out over time to forecast foreign exchange rate using different machine learning methods. This literature review examines key theoretical frameworks, empirical findings, and emerging methodologies in foreign exchange rate forecasting, identifying trends and gaps in the existing research.

Authors in [6] built an effective model for predicting forex price trends by leveraging Recurrent Neural Networks (RNNs), with a particular focus on LSTM networks. This research utilized secondary data comprising historical forex prices from several financial markets over an extended period. The data was meticulously preprocessed and divided into a training set and a testing set, following the standard approach of assigning 70% to training and 30% to testing. The performance of the model was evaluated using Mean Squared Error (MSE) and Mean Absolute Error (MAE) as key metrics. The LSTM model achieved an MSE value of 0.003052 and a MAE of 0.002390. Based on these findings, the study recommended the use of LSTM networks for forex price trend forecasting in financial institutions and trading firms as LSTM is superior to traditional methods for forecasting forex price trends. Additionally, the comparative analysis of traditional statistical methods with modern machine learning techniques provided a comprehensive evaluation of the models' performance, highlighting the advantages of using RNNs for time series forecasting. The study also contributed to the growing body of literature on the application of deep learning techniques in finance, demonstrating the practical benefits of these methods in enhancing prediction accuracy. The detailed analysis and comparison of various models ensured that the study's conclusions were well-rounded and supported by empirical evidence.

Yildirim et al. explored the use of LSTM models for forecasting the directional movement of the EUR/USD currency pair over different time horizons, specifically one day, three days, and five days ahead. The study introduced a novel performance metric, profit_accuracy, to evaluate the effectiveness of the predictions in generating profitable transactions. The data included both macroeconomic indicators and technical indicators, which were used to train and evaluate the LSTM models. The research applied two separate LSTM models, one trained using macroeconomic data (ME_LSTM) and the other using technical indicators (TI_LSTM). A classifier was developed to determine the directional movement of the EUR/USD pair into three classes: no_action, decrease, and increase. The hybrid model, which combined both macroeconomic and technical indicator features (ME_TI_LSTM), was also tested. The study found that the ME_LSTM model slightly outperformed the TI_LSTM model in terms of both profit_accuracy and the number of transactions generated, though the difference was minimal and statistically insignificant. The hybrid model (ME_TI_LSTM), which combined all features, did not show a significant improvement in accuracy compared to the individual models. However, the proposed hybrid model demonstrated the best overall performance, achieving an average profit_accuracy of 73.61% across all prediction periods. It also reduced the number of transactions by 40.37% on average compared to the baseline models, primarily by dropping risky transactions. In conclusion, the study provided compelling evidence that the hybrid LSTM model offers a robust approach to forecasting the directional movement of the EUR/USD currency pair [15].

Authors in [12] explored the effectiveness of LSTM networks, particularly when combined with event-driven inputs, for predicting foreign exchange rates. The study focused on developing a model that could leverage both the sequential patterns inherent in time series data and the impact of specific events, which often cause significant fluctuations in forex markets. The data utilized in this research consisted of historical forex rates and event data over a substantial period, including major economic announcements, geopolitical events, and other news that could influence forex prices. The LSTM model was constructed, incorporating event-driven features to enhance its predictive capabilities. The hybrid approach aimed to address the limitations of traditional LSTM models, which may not fully capture the impact of sporadic, yet significant events on forex prices. The study's results demonstrated that the event-driven LSTM model significantly improved the accuracy of forex price predictions compared to standard LSTM models and other traditional methods. The performance of the model was evaluated using metrics such as Root Mean Error (RME), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). The event-driven LSTM achieved an RME of 0.006*10^(-3), RMSE of 2.407*10^(-3), MAE of 1.708*10^(-3), and MAPE of 0.194% highlighting its effectiveness in capturing the complex, event-driven patterns within the forex market.

Authors in [11] explored the effectiveness of combining Perceptron with Genetic Algorithms (GAs) for predicting foreign exchange rates. The study focused on developing a hybrid model that leverages the strengths of both techniques. The data utilized in this research consisted of historical forex rates from multiple currency pairs over a substantial period. The hybrid model aimed to enhance prediction accuracy by addressing the limitations of individual methods. The results demonstrated that the hybrid Perceptron-GA model significantly improved the accuracy of forex rate predictions compared to traditional methods and standalone Perceptron models. The evaluation of model performance was based on metrics such as Mean Squared Error (MSE) and Mean Absolute Error (MAE). The hybrid model achieved an MSE of 0.01 and a MAE of 0.0082. This result underscored the effectiveness of integrating Perceptron and GAs in capturing the complex, non-linear patterns inherent in forex data. One of the key strengths of this study was its innovative approach to combining machine learning and evolutionary algorithms to tackle a real-world financial problem. The research also contributed to the broader field of financial forecasting by demonstrating the potential of hybrid models in enhancing prediction accuracy. The focus on optimizing neural networks using evolutionary algorithms added a new dimension to the existing literature, showcasing the practical benefits of this approach in finance.

Authors in [14] investigated the effectiveness of tree ensemble methods, including Random Forests, Gradient Boosting Machines (GBMs), and Extreme Gradient Boosting (XGBoost), for predicting trends in the forex market. The study utilized historical forex rate data spanning multiple currency pairs over a significant time period. The results of the study indicated that XGBoost significantly outperformed traditional methods and simpler machine learning models in predicting forex market trends. The performance of the models was evaluated using metrics such as Mean Squared Error (MSE) and Mean Absolute Error (MAE), with XGBoost achieving MSE of 0.009 and MAE of 0.0075. These results underscored the models' ability to handle the intricacies of forex data, making them highly effective for trend prediction. The research also contributed to the broader field of financial forecasting by demonstrating the effectiveness of tree ensemble methods in capturing complex patterns in forex data.

Despite the importance of exchange rate fluctuations for Nepal's economy, little research has focused on forecasting forex rates in the Nepalese context. Most studies target developed countries, with limited use of advanced machine learning methods. This study aims to fill that gap by applying and comparing GRU and LSTM models to predict the Nepali Rupee exchange rate, identifying which model offers better forecasting accuracy.

3. Methodology

Forecasting foreign exchange rates involves predicting the value of one currency relative to another. Accurate predictions are crucial for various stakeholders, including investors, businesses engaged in international trade, and policymakers. Effective forecasting helps in managing currency risk, optimizing investment strategies, and implementing sound economic policies [3]. In Nepal, which has a significant reliance on remittances and international trade, precise forex rate forecasts are vital for maintaining economic stability and improving growth. The analysis of forecasting forex rates in Nepal using LSTM and GRU models is rooted in the intersection of financial forecasting and advanced machine learning techniques. Accurate prediction of forex rates is essential for effective financial decision-making, economic planning, and risk management. This study focuses on applying two advanced machine learning models: LSTM and GRU to forecast the forex rate of Nepal in relation to major currency such as USD.

3.1. Data Collection

The dataset used in this study consists of time series data encompassing historical records relevant to forex rates. This data is sourced from secondary sources, specifically from the Nepal Rastra Bank. The time series dataset includes historical forex rate information from 2005 to 2024 that provides insights into past currency fluctuations, which is essential for evaluating the forecasting capabilities of the LSTM and GRU models. This historical data provides the foundation for evaluating the forecasting performance of the LSTM and GRU models. The data includes time series records of forex rates between the Nepalese Rupee and major currency, the US Dollar.

3.2. Stationarity Test

For an accurate forecasting of a time series forecasting, a key concept is stationarity. Stationarity in the context of time series data refers to the property of a time series where its statistical properties, such as mean, variance, and autocorrelation, remain constant. For time series data to be stationarity, it should not exhibit trends, seasonality, or other time-dependent structures that cause its statistical properties to change over time. When the data is stationary, it is easier to model and forecast future values, as the past behavior of the time series can be used to predict future behavior. The Augmented Dickey-Fuller test is a parametric method used to assess whether a unit root is present in a dataset. The existence of a unit root suggests that the data is non-stationary, indicating that it may display a trend or seasonal pattern. To perform the ADF test, an autoregressive model with a differencing term is fitted to the data, and the significance of the differencing coefficient is evaluated.

Table 1. Augmented Dickey-Fuller Test

Variables	**Dickey-Fuller**	**Lag order**	**p-value**	**Stationarity**
Actual Forex Rate	-2.9564	0	0.1732(>0.05)	Non-stationary
1^{st} order difference	-16.0064	1	0.0100(<0.05)	Stationary

The table 1 summarizes the results of the ADF test applied to the actual forex rate and its first-order difference to determine stationarity. For the actual forex rate, the Dickey-Fuller test statistic is -2.95645431195, with a lag order of 0. The corresponding p-value is 0.173228, which is greater than the significance level of 0.05. This result indicates that the null hypothesis of a unit root cannot be rejected, suggesting that the actual forex rate series is non-stationary. In contrast, when the first-order difference of the forex rate is tested, the Dickey-Fuller test statistic is -16.0064402294 with a lag order of 1. The p-value for this test is 0.010, which is significantly less than 0.05. This implies that the null hypothesis of a unit root is rejected, indicating that the first-order differenced series is stationary. Therefore, while the original forex rate series exhibits non-stationarity, differencing the series once renders it stationary, which is essential for further time series analysis and forecasting.

3.3. Data Preprocessing and Splitting

The collected data was cleaned to address any inconsistencies or missing values. This included handling outliers, correcting data entry errors, and interpolating missing values where necessary to ensure a complete and accurate dataset. The data was then transformed to make it suitable for analysis using normalization and feature engineering. Normalization was done for scaling the data to a uniform range to improve the performance and convergence of the models. Min-Max scaling technique was applied for normalization. Feature engineering was done for creating relevant features from the raw data that can enhance the predictive capability of the models. This includes lag variables, and moving averages. The dataset was divided into training and test subsets using an 80-20 split ratio. This division allows for robust model training and evaluation, helping to assess the model's performance on unseen data.

3.4. Model Description

Long Short-Term Memory (LSTM) is a type of recurrent neural network (RNN) designed to address the limitations of traditional RNNs, particularly the vanishing gradient problem. Proposed by [9], LSTMs are highly effective in capturing long-term dependencies in sequential data, making them well-suited for time series forecasting tasks such as foreign exchange rate prediction. Unlike standard RNNs, LSTMs incorporate specialized gating mechanisms – the forget gate, input gate, and output gate – which regulate the flow of information through the network. These gates enable LSTMs to retain relevant historical data while discarding irrelevant information, leading to improved predictive performance in highly volatile financial markets. Due to their ability to model complex temporal relationships, LSTMs have been widely applied in financial time series forecasting, demonstrating superior accuracy compared to traditional econometric models like ARIMA and GARCH [7]. By capturing both short-term and long-term dependencies, LSTMs are highly effective for tasks like time series forecasting, where the future values depend on a complex and non-linear relationship with past observations [8]. The different equations of LSTM are as follows.

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

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

$\tilde{C}_t = tanh(W_C \,.\, [h_{t-1}, x_t] + b_C)$ Candidate cell state

$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t$ Update cell state

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

$h_t = o_t * \tanh(C_t)$ Hidden state update

GRU is a variant of recurrent neural networks (RNNs) introduced by [4] to address the vanishing gradient problem and improve sequence modeling efficiency. Similar to LSTM, GRU is designed to capture long-term dependencies in sequential data, making them suitable for time series forecasting tasks such as foreign exchange rate prediction. However, GRUs have a simpler architecture than LSTMs, as they use only two gates – the reset gate and update gate – instead of three. This reduction in complexity allows

GRUs to achieve comparable performance with LSTMs while being more computationally efficient [5]. Recent studies have demonstrated the effectiveness of GRU-based models in financial forecasting, showing their ability to adapt to the nonlinear and volatile nature of exchange rate movements. Given their efficiency and accuracy, GRUs have become a popular alternative to LSTMs in deep learning applications for time series analysis. The different equations of GRU are as follows.

$r_t = \sigma(W_r \, . [h_{t-1}, x_t] + b_r)$ Reset gate

$z_t = \sigma(W_z \, . [h_{t-1}, x_t] + b_z)$ Update gate

$\tilde{h}_t = tanh(W_h \, . [r_t * h_{t-1}, x_t] + b_h)$ Candidate hidden state

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

3.5. Experimental Setup

Both LSTM and GRU algorithms were implemented using Python Programming Language and the libraries such as Keras, Pandas, NumPy, and Matplotlib. The models underwent training and testing on Microsoft Windows 11, AMD Ryzen 5 5500U CPU @ 2.10 GHz, and 8 GB RAM.

3.6. Performance Evaluation

The performance of both LSTM and GRU models are assessed using Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). MSE quantifies the average squared difference between the predicted values and the actual values. It helps to evaluate how well a model performs. A lower MSE indicates a better model.

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \qquad (1)$$

RMSE measures the average magnitude of errors by taking the square root of the MSE. This metric expresses the error in the same units as the predicted values, which makes it easier to interpret compared to MSE. It penalizes large errors more than small ones. Lower RMSE values indicate that the model's predictions are closer to the actual values, and the model has fewer and smaller errors overall.

$$RMSE = \sqrt{MSE} \qquad (2)$$

MAE measures the average absolute difference between actual and predicted values. It provides a straightforward interpretation of how far predictions are from actual values. It is less sensitive to outliers compared to MSE and RMSE. It expresses the error in the same unit as the original data, making it easy to interpret. A lower MAE indicates a more accurate model.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \qquad (3)$$

MAPE evaluates the accuracy of the model's predictions as a percentage of the actual values. This metric provides a relative measure of error, making it useful for comparing the performance of different forecasting models or datasets. MAPE calculates the average percentage by which the predicted values deviate from the actual values. A lower MAPE indicates that the model's forecasts are more accurate in percentage terms, giving a clear understanding of prediction accuracy relative to the size of the actual values.

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right| \times 100 \qquad (4)$$

4. Result Analysis

The Figure 1 presents the historical trend of the Nepalese rupee against the US dollar over a period of days. The chart reveals a steady upward movement in the forex rate, showing a gradual appreciation of the Nepalese rupee against the US dollar. A notable point in the chart is the period of increased volatility, which occurs around the earlier days, particularly between days 3000 and 3500. This fluctuation suggests that there were significant market events or interventions influencing the exchange rate during that time. From day 3500 onward, the chart shows a consistent increase in the forex rate with only occasional small fluctuations. The general upward trend implies a gradual depreciation of the Nepalese rupee, indicating that more units of the rupee are needed to purchase one US dollar as time progresses.

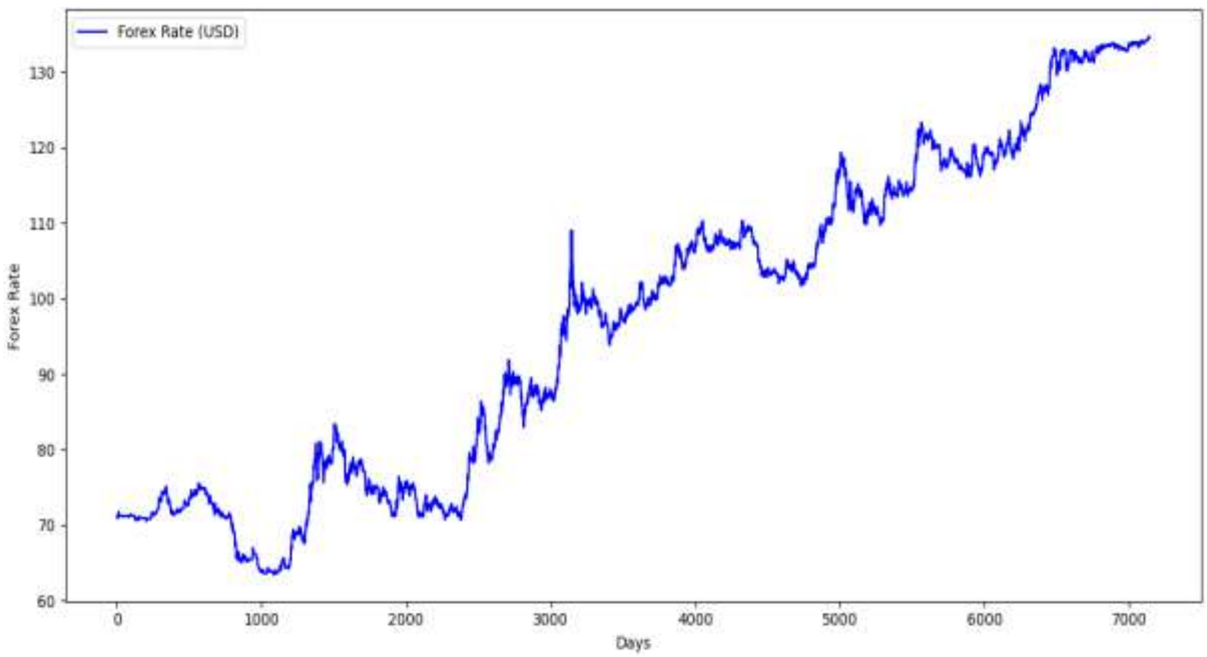

Figure 1. Trendline of Forex Rate of Nepal (2005-2024)

4.1. Analysis of LSTM

The Figure 2 illustrates the training loss of LSTM model over 100 epochs. The vertical axis represents the loss value, which quantifies the difference between the predicted and actual values during training, while the horizontal axis shows the number of epochs. The chart indicates a rapid decrease in the training loss within the first few epochs, with the loss value dropping sharply from above 0.52 to around 0.48 within the first 5 epochs. After this initial decline, the training loss plateaus and stabilizes, maintaining a relatively constant value just below 0.48 for the remainder of the training process up to 100 epochs.

The pattern shows that the model learns quickly during the initial training phase, with a steep loss decline. After that, further training offers minimal improvement, suggesting diminishing returns. This stabilization indicates the model has minimized error, and additional training may lead to overfitting without significant gains.

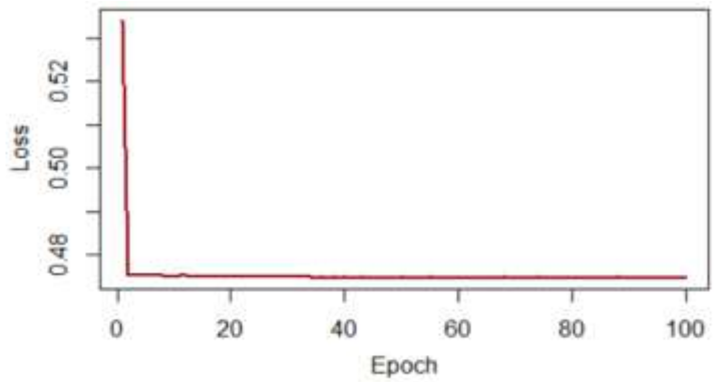

Figure 2. Training Loss of LSTM

Table 2. Summary Result of LSTM Model

Input Feature	**Metrics**			
Unit	**MSE**	**RMSE**	**MAE**	**MAPE**
30	7.0025	2.6462	2.4719	1.9225%
40	5.3176	2.3059	2.1393	1.6623%
50	4.7056	2.1692	2.0262	1.5764%
60	4.9369	2.2219	2.1189	1.6547%

The Table 2 provides a detailed comparison of the performance metrics for testing dataset of LSTM with varying numbers of unit while maintaining the same epoch count of 50, batch size of 64, and dropout rate of 0.3. The models differ only in the number of LSTM units, and their performance is evaluated using four metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE).

For the LSTM model with 30 units, the MSE is 7.0025, indicating the average squared difference between actual and predicted values. The RMSE, which translates this into the same units as the data, is 2.6462, providing a tangible measure of the prediction error. The MAE is 2.4719, representing the average absolute magnitude of prediction errors, and the MAPE is 1.9225%, suggesting a relatively low prediction error as a percentage of the actual values. When the number of units is increased to 40, the model shows improved performance across all metrics. The MSE decreases to 5.3176, and the RMSE to 2.3059, reflecting a reduced prediction error. Similarly, the MAE drops to 2.1393, and the MAPE to 1.6623%, indicating better accuracy and robustness in predictions. The model with 50 units continues this trend, achieving the lowest overall MSE of 4.7056, RMSE of 2.1692, MAE of 2.0262, and MAPE of 1.5764% among all models. These metrics suggest this model provides the most accurate and reliable predictions within this setup. However, increasing the number of units to 60 slightly increases the errors, with an MSE of 4.9369, RMSE of 2.2219, MAE of 2.1189, and MAPE of 1.6547%. While still competitive, this model demonstrates that adding more units does not always lead to better performance, potentially due to diminishing returns in predictive accuracy.

Among all the models evaluated, the LSTM model with 50 units demonstrates the best overall performance, achieving the lowest error metrics. This indicates that the configuration with 50 units is optimal for balancing complexity and predictive accuracy for the forecasting of forex rate of Nepal.

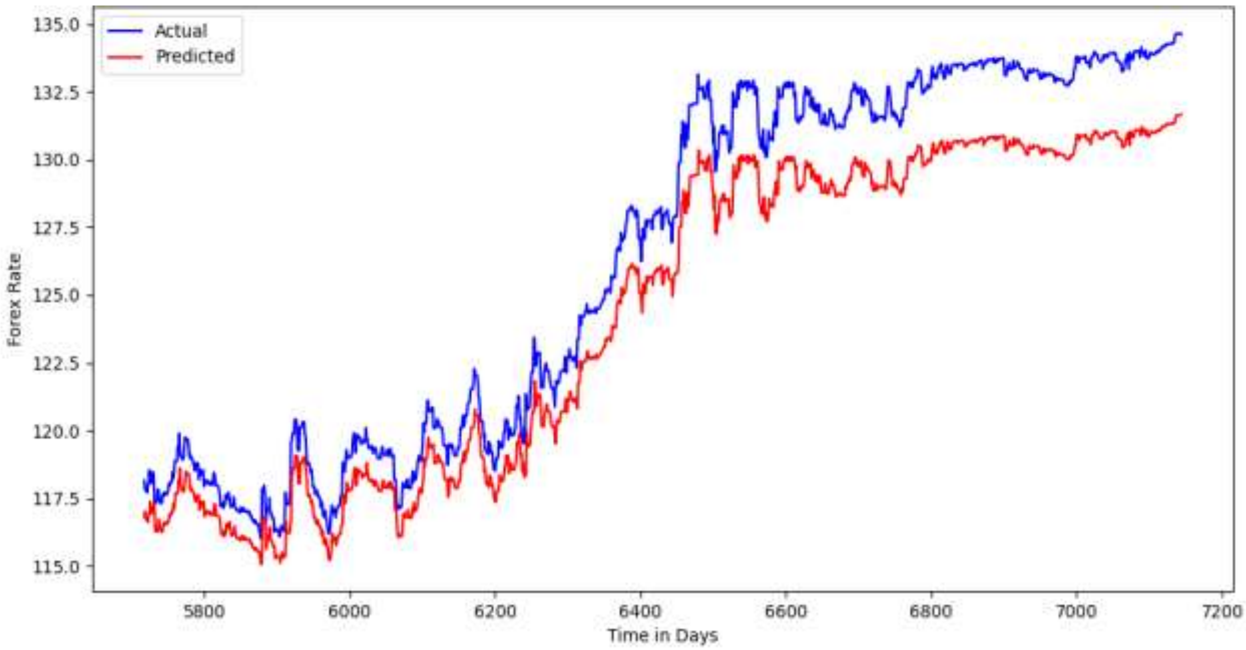

Figure 3. Actual Vs Predicted Forex Rate of LSTM

The Figure 3 illustrates the performance of the testing dataset for the best-performing Long Short-Term Memory (LSTM) model, identified as the LSTM model trained with 50 units. The blue line represents the actual forex rates, while the red line depicts the predicted forex rates generated by the model. From the figure, it is evident that the model closely follows the actual forex rates over time, particularly in the earlier stages, where the blue and red lines are nearly overlapping. This close alignment demonstrates the model's ability to effectively learn and replicate the trends and patterns present in the historical data.

As the time progresses, however, there is a noticeable divergence between the actual and predicted values, particularly in the latter part of the dataset. This gap suggests that while the model is highly accurate, it may face challenges in fully capturing the more volatile or less predictable fluctuations in forex rates in these later periods.

Overall, the small differences between the actual and predicted values align with the previously discussed performance metrics, such as the low Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). These metrics underscore the model's robustness and accuracy in forecasting forex rates, making it a reliable tool for time series predictions in scenarios requiring precision, such as currency exchange rate forecasting.

4.2. Analysis of GRU

The Figure 4 illustrates the training loss of an GRU model as a function of the number of epochs during the training process. The vertical axis represents the loss, while the horizontal axis shows the number of epochs, ranging from 0 to 70. The chart indicates that the training loss starts relatively high at the beginning of the training. However, it rapidly decreases within the first few epochs, which is a common pattern indicating that the model is quickly learning the key patterns in the data. After about 10 to 20 epochs, the loss begins to stabilize, approaching a value close to zero.

This flattening of the curve suggests that the model has reached a point where further training brings minimal improvements in reducing the loss, indicating convergence. The steady loss value towards the end of the epochs indicates that the model has effectively learned from the training data, and additional epochs are unlikely to significantly improve its performance. This suggests that the training process is efficient and that the chosen number of epochs is sufficient for the model to achieve optimal performance without overfitting.

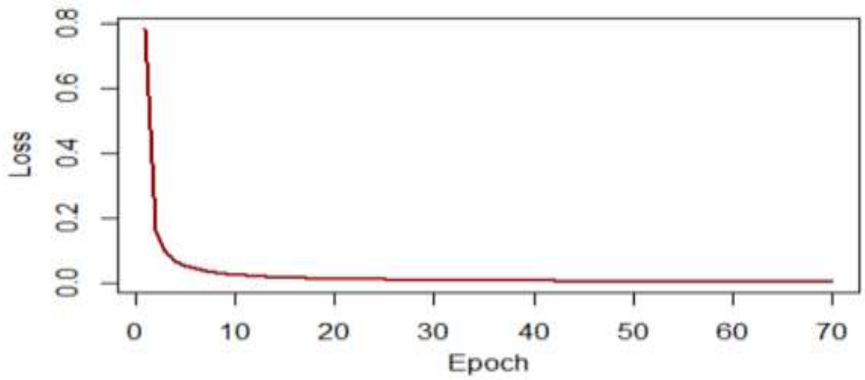

Figure 4. Training Loss of GRU

Table 3. Summary Result of GRU Model

Input Feature	Metrics			
Unit	**MSE**	**RMSE**	**MAE**	**MAPE**
30	20.4508	4.5222	4.4064	3.4539%
40	11.5716	3.4017	3.2905	2.5753%
50	7.1607	2.6759	2.5673	2.0065%
60	8.3174	2.8839	2.8074	2.2012%

The Table 3 provides a detailed comparison of the performance metrics for the testing dataset of Gated Recurrent Unit (GRU) models with varying numbers of units, while maintaining a constant epoch count of 50, batch size of 64, and dropout rate of 0.3. The models are evaluated using four metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE).

For the GRU model with 30 units, the MSE is 20.4508, which is relatively high compared to the other configurations. This translates into an RMSE of 4.5222, indicating a significant prediction error in the same units as the data. The MAE is 4.4064, reflecting the average magnitude of prediction errors, while the MAPE of 3.4539% suggests a moderate percentage error relative to the actual values. Increasing the number of units to 40 leads to substantial improvements in performance. The MSE decreases to 11.5716, and the RMSE to 3.4017, showing a notable reduction in prediction error. Similarly, the MAE drops to 3.2905, and the MAPE to 2.5753%, indicating enhanced accuracy and better trend-capturing capability. The model with 50 units achieves the best overall performance, with the lowest MSE of 7.1607, RMSE of 2.6759, MAE of 2.5673, and MAPE of 2.0065%. These metrics confirm that this configuration provides the most accurate and reliable predictions for the testing dataset. However, increasing the number of units to 60 slightly worsens the performance, as evidenced by an MSE of 8.3174, RMSE of 2.8839, MAE of 2.8074, and MAPE of 2.2012%. This suggests that adding more units beyond a certain point may lead to diminishing returns, possibly due to an increase in model complexity without corresponding gains in predictive accuracy.

Among the GRU models evaluated, the configuration with 50 units stands out as the optimal choice, achieving the lowest error metrics across all categories. This indicates that the model with 50 units strikes the best balance between complexity and accuracy, making it the most effective for forecasting tasks in this scenario.

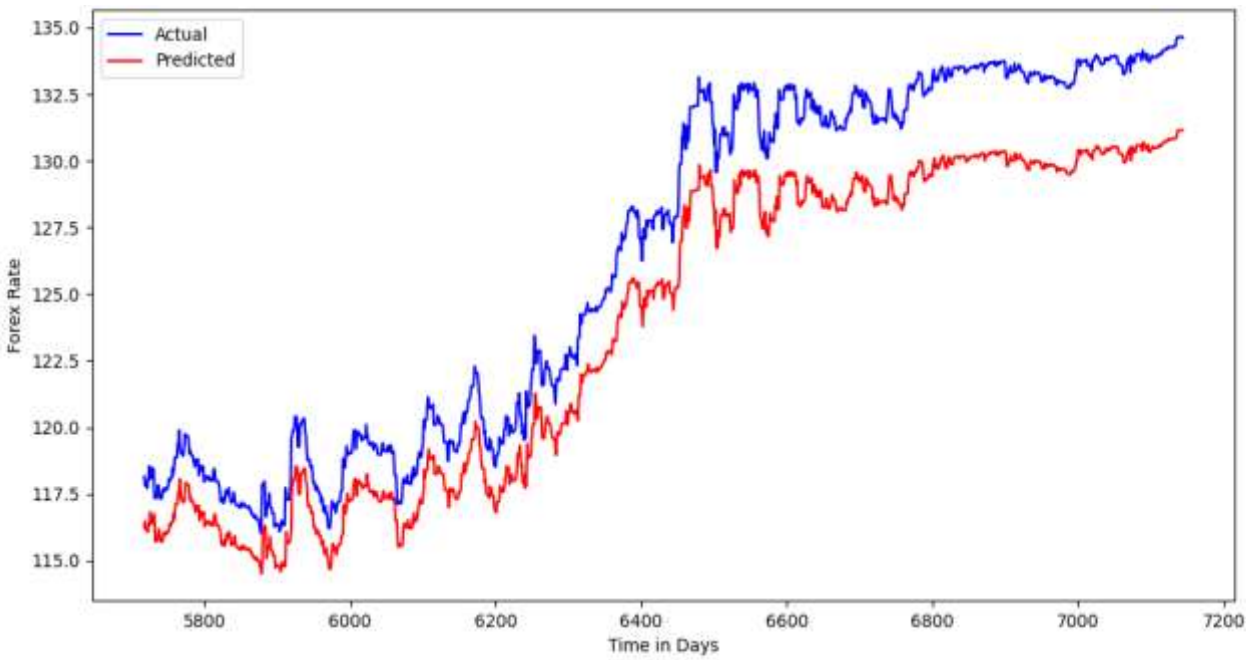

Figure 5. Actual Vs Predicted Forex Rate of GRU

The Figure 5 illustrates the performance of the testing dataset for the best-performing GRU model, identified as the GRU model trained with 50 units. The blue line represents the actual forex rates, while the red line depicts the predicted forex rates generated by the model. From the figure, it is clear that the model closely tracks the actual forex rates over time, especially in the earlier segments where the blue and red lines almost overlap. This alignment highlights the model's capability to accurately capture the underlying trends and patterns in the historical data, showcasing its predictive strength.

However, as time progresses, a slight divergence between the actual and predicted values becomes apparent, particularly in the latter stages of the dataset. This gap indicates that while the model performs exceptionally well overall, it encounters some difficulty in fully capturing the more volatile or less predictable fluctuations in forex rates toward the end of the time series.

The small differences between the actual and predicted values reflect the previously discussed performance metrics, such as the low Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). These metrics confirm the model's robustness and accuracy in forecasting forex rates, making it a reliable tool for time series predictions in practical applications like currency exchange rate forecasting.

4.3. Model Comparison

This research focuses on identifying the most accurate and reliable model for predicting forex rates. To achieve this, best-performing versions of each Long Short-Term Memory and Gated Recurrent Unit models are selected for comparison. These versions are chosen based on their superior performance metrics, such as Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). By focusing on these top-performing models, the comparison aims to determine which architecture: LSTM or GRU offers the most effective and precise predictions for the given data and task.

Table 4. Model Comparison

Model	**Metrics**			
	MSE	**RMSE**	**MAE**	**MAPE**
LSTM	4.7056	2.1692	2.0262	1.5764%
GRU	7.1607	2.6759	2.5673	2.0065%

The Table 4 presents a comparative analysis of two models, LSTM and GRU, using four evaluation metrics: Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). The MSE value of the LSTM model is 4.7056, which indicates that average squared difference between the predicted and actual values, compared to MSE of 7.1607 for GRU model. The RMSE value of the LSTM model is 2.1692, which measures the squared root of average squared difference between predicted values, compared to RMSE of 2.6759 for GRU. The average absolute difference between the predicted and actual values of LSTM model is 2.0262 compared to that of 2.5673 of GRU model. The MAPE value of 1.5764% of LSTM model means that the average difference between the forecasted value and the actual value is 1.5764% compared to that of 2.0065% of GRU model. Lower values of MSE, RMSE, MAE, and MAPE indicate better model performance. Thus, LSTM model outperforms GRU model in all four of the metrics for forecasting the forex rate of Nepal.

5. Conclusion

This study employed historical data of forex rate of Nepali Rupee against US dollar. The Augmented Dickey-Fuller (ADF) test was employed to assess stationarity. The result showed that the actual forex rate was non-stationary (p-value = 0.173228 > 0.05). However, after first-order differencing, the series becomes stationary (p-value = 0.010 < 0.05). Based on the ADF test results, a lag of 1 was chosen for model computations. The dataset was split into training (80%) and testing (20%) subsets. Min-Max scaling was applied to normalize the features, ensuring all features contribute equally and improving the convergence of learning algorithms. Both LSTM and GRU models were implemented using Python programming language. The best performance was achieved by the LSTM model with 50 units, showing the MSE of 4.7056, RMSE of 2.1692, MAE of 2.0262, and MAPE of 1.5764%. Similarly, the best performance was observed in the GRU model with 50 units, with the lowest MSE of 7.1607, RMSE of 2.6759, MAE of 2.5673, and MAPE of 2.0065%. The analysis demonstrated that the LSTM model significantly outperformed the GRU model in accuracy and precision, as shown by lower MSE, RMSE, MAE, and MAPE metrics. The LSTM model's superior ability to capture historical trends and handle time series data with minimal error highlighted its robustness in forecasting volatile exchange rates. The findings indicated that the LSTM model is more effective in predicting the forex rate of Nepal. Given the ability of LSTM model to effectively capture historical trends and handle time series data with minimal error, it is recommended to use LSTM models for forecasting the forex rate of Nepal.

Although the study focused on historical forex rates of NPR against US Dollar, incorporating additional economic indicators such as interest rates, Foreign Direct Investment (FDI) etc. could further enhance the model’s predictive capabilities. The collaboration between academic institutions, financial organizations, and government bodies should be encouraged to promote the use of advanced forecasting techniques for forex forecasting.

References

[1] 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(9):**.

[2] Box, G. E. P., Jenkins, G. M., & Reinsel, G. C. **(2008)**. Time Series Analysis: Forecasting and Control. In *Time Series Analysis: Forecasting and Control,* (4th ed.). Wiley.

[3] Cheung, Y. W., & Chinn, M. D. **(2001)**. Currency traders and exchange rate dynamics: A survey of the US market. *Journal of International Money and Finance*, **20(4)**.

[4] Cho, K., van Merrienboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. **(2014)**. *Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation*.

[5] Chung, J., Gulcehre, C., Cho, K., & Bengio, Y. **(2014)**. *Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling*.

[6] Dobrovolny, M., Soukal, I., Lim, K. C., Selamat, A., & Krejcar, O. **(2020)**. Forecasting of FOREX Price Trend using Recurrent Neural Network - Long short-term memory. *Proceedings of the International Scientific Conference Hradec Economic Days 2020*, **10**.

[7] Fischer, T., & Krauss, C. **(2018)**. Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, **270(2)**: 654–669.

[8] Greff, K., Srivastava, R. K., Koutnik, J., Steunebrink, B. R., & Schmidhuber, J. **(2017)**. LSTM: A Search Space Odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, **28(10)**.

[9] Hochreiter, S., & Schmidhuber, J. **(1997)**. Long Short-Term Memory. *Neural Computation*, ***9*(8)**.

[10] 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)**: 3–24.

[11] Meirinaldi, Yolanda, & Seputra, Y. E. A. **(2022)**. Analysis of Foreign Exchange Using Perceptron and Genetic Algorithm Machine Learning (GALM). *European Journal of Business and Management Research*, **7(6)**.

[12] Qi, L., Khushi, M., & Poon, J. **(2020)**. Event-Driven LSTM for Forex Price Prediction. *2020 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*.

[13] Shrestha, M. B., & Chaudhary, R. K. **(2012)**. The impact of exchange rate fluctuations on exports: A case study of Nepal. *Journal of International Economics and Economic Policy*, **9(1)**: 79–90.

[14] Thao, N., Nguyen, H.-C., Mach, B.-N., Thuan, D., Quynh, T., Huong, T., Chi, D., & Nguyen, T. **(2024)**. *Unlocking Forex Market Trends: Advanced Predictive Modeling with Tree Ensembles*.

[15] Yıldırım, D. C., Toroslu, I. H., & Fiore, U. **(2021)**. Forecasting directional movement of Forex data using LSTM with technical and macroeconomic indicators. *Financial Innovation*, **7(1)**.

□□