EAI Endorsed Transactions
on Tourism, Technology and Intelligence
Research Article EAI.EU

A Hybrid Ensemble Deep Learning and Reservoir Computing Approach for Exchange Rate Prediction

Senthan Prasanth* and Chau Bao Phung

Faculty of Engineering and Applied Science, Memorial University, Canada

Abstract

A nation's currency plays a vital role in stabilizing the economy and determining exchange rates in global markets. Keeping track of the influential currencies while comparing them with one's own currency becomes essential these days. During this study, we have specifically focused on the exchange rate prediction between the United States dollar (USD) and the Canadian dollar (CAD). This pair is one of the most active currency pairs with significant economic implications for both nations. This paper studies the use of machine learning models for this specific matter, which includes long short-term memory (LSTM), gated recurrent units (GRU), and reservoir computing (RC) echo state network models. They were evaluated not only as individual models but also in various hybrid combinations. A hybrid model which combines LSTM and RC yielded better performance in terms of Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Coefficient of Determination (R^2). Moreover, the full potential of RC represents a promising direction for future researchers to incorporate into time series analysis. In addition, considering internal and external factors that influence the exchange rate during model development would also give more accurate predictions.

Received on 30 September 2025; accepted on 18 December 2025; published on 22 December 2025

Keywords: Exchange rate forecasting, Reservoir computing, Echo state networks, Time series

doi:10.4108/eettti.10443

1. Introduction

An exchange rate represents the value of a nation's currency when it is traded for another currency. The relative strength or weakness of a currency has significant implications for a wide range of economic activities such as international trade [1], [2], [3], investment decisions [4], and risk management [4], [5]. However, exchange rate prediction is widely regarded as one of the most complex challenges in financial forecasting. This complexity arises due to non-linearity [1], [6], [7], and non-stationarity [1], [7], [8] of foreign exchange markets. Although in many countries, it is controlled by the government through various policies, the exchange rate is still influenced by certain external factors, namely, geopolitical events [2], [4], [9], economic indicators [1], [3], and market sentiment [2], [3].

This study specifically focuses on the United States dollar (USD) to Canadian dollar (CAD) exchange rate due to high volatility in recent times [3], [9]. The USD/CAD pair, often referred to as the 'Loonie' [10], is influenced by factors such as oil prices [3], interest rate differentials between the Federal Reserve and Bank of Canada [3], [9], and economic disparities between the United States and Canada [3].

Traditional statistical models [8], [4], [11] have struggled to capture exchange rates due to the aforementioned complexities. These limitations have triggered researchers to explore advanced computational approaches, especially in the area of artificial intelligence. Moreover, deep learning (DL) methods like long short-term memory (LSTM) and gated recurrent units (GRU) have shown prominent results in financial time series prediction [5], [12]. These models also produced promising results in modelling temporal dependencies. At the same time, reservoir computing (RC) [13], [14], [15] has grabbed attention for its computational efficiency and ability to model chaotic systems with minimal training requirements [14], [15]. While individual models have their strengths, they also exhibit

*Corresponding author. Email: sprasanth@mun.ca

limitations when applied to the multifaceted challenge of exchange rate forecasting.

This research makes three primary contributions: first, evaluating the individual performance of LSTM, GRU, and RC models for USD/CAD exchange rate prediction; second, developing and comparing various hybrid ensemble combinations of these architectures; and third, exploring the computational efficiency and predictive capabilities of reservoir computing when integrated into hybrid frameworks.

The rest of the paper is organized as follows: Section 2 details a comprehensive literature review of existing approaches in exchange rate forecasting. Section 3 outlines the data, methodology, and model configurations. Moreover, section 4 presents our experimental findings and the obtained results. The final section concludes the paper with future work to carry out.

2. Traditional Statistical Methods and ML/DL Approaches in Exchange Rate Prediction and Application of RC in Financial Forecasting

Exchange rate prediction has been researched for a long period using several methodologies. Early approaches primarily focused on statistical models [8] before the emergence of ML and DL techniques [5]. This section explores both traditional methods and ML/DL approaches used in exchange rate prediction. Moreover, since RC has not been specifically applied to exchange rate prediction, its successful implementation in the financial domain is discussed.

2.1. Statistical Methods and ML/DL Approaches for Exchange Rate Prediction

Econometric and statistical models, such as the autoregressive moving average (ARMA) [11], autoregressive fractionally integrated moving average (ARFIMA), and autoregressive integrated moving average (ARIMA) [4], [11], have been commonly used for exchange rate forecasting.

Henri Makika and her teammates [4] tested neural network models such as Multilayer Perceptron (MLP), LSTM, GRU, and a hybrid model called AE-BiGRU-Swish for forecasting the exchange rates. Moreover, they compared the performance of the aforementioned models with traditional approaches [4]. Moreover, they collected Brazilian Real to US Dollar and Euro to US Dollar rates for the study. From the observation, it was explored that the proposed AE-BiGRU-Swish model had the best overall performance in predicting the Brazilian dollar, particularly for shorter forecasting horizons (1-7 days). A similar kind of superior performance results were obtained for forecasting Euro as well [4].

Research by Emilio Colombo and Matteo Pelagatti [8] proved that statistical learning methods can improve exchange rate forecasting compared to traditional models. They carried out the research considering 18 currencies, which were evaluated against the USD. Moreover, they tested the performance of those 18 currencies' exchange rates on five models, namely standard regression (REG), elastic net on splines (SP), elastic net on splines with interactions (SPI), random forest (RF), and support vector machines (SVM). Based on the findings, it was clearly noted that SVM outperformed all the other models [8].

Yet in another research, Deep et al. [11] conducted a study on predicting the Australian Dollar (AUD) against the US Dollar using the ARIMA model enhanced with economic indicators. Their study analyzed daily foreign exchange rates from January 2016 to December 2020 and predicted the 2021 exchange rate with high precision. They achieved an impressive 96.98% forecasting accuracy with their optimized model [11].

A hybrid model combining GRU and LSTM was developed [5] for predicting the closing prices of the Euro to USD, the Great British Pound to USD, USD/CAD, and USD to Swiss Franc. The architecture consists of a GRU layer followed by an LSTM layer. This hybrid model aims to utilize the speed of GRU and the accuracy of LSTM with longer sequences. The study revealed that exchange rates predicted by the GRU-LSTM hybrid model significantly outperformed those predicted by standalone LSTM and GRU models. It also maintained superiority in terms of risk associated with return compared to a simple statistical model. However, sometimes, it was noticed that this model also struggles with sudden price changes [5].

Linkon and his team [16] also worked on exchange rate prediction, where they researched to forecast the Japanese Yen against the USD. In their research, they implemented models like LSTM, GRU, Generative Adversarial Networks (GAN), and Wasserstein Generative Adversarial Networks (WGAN). The WGAN model excelled in generalization, showing the best predictive performance (R-Square: 0.9785) on unseen test data. While GAN performed marginally better on the training data, WGAN proved more robust when forecasting actual future exchange rates [16].

2.2. Applications of Reservoir Computing in Financial Forecasting

As traditional statistical models often fail to adequately represent the aforementioned complexities [8], [4], [11], reservoir computing (RC), particularly the echo state network (ESN), has emerged as a promising approach that balances computational efficiency with predictive power [13], [14], [15].

Wang and her team [14] applied the ESN to predict seven international stock market indices. They compared the performance of ESN with that of two other techniques, namely LSTM and recurrent neural network (RNN), for all seven indices. In addition, to further validate their approach, they evaluated the performance of two other models, namely EMD2FNN and ModAugNet. They tested EMD2FNN using three indices (S&P500, NASDAQ, and SSE) and ModAugNet using two indices (S&P500, KOSPI200) against ESN. For the S&P500 index, the ESN slightly underperformed compared to EMD2FNN, but the results were almost alike. However, for NASDAQ and SSE indices, the ESN achieved significant improvements over EMD2FNN. Similarly, when compared against ModAugNet, the ESN outperformed on S&P500 prediction with 16.63% and 19.44% decreases in MSE and MAPE, respectively. But it slightly underperformed on KOSPI200 [14].

A novel hybrid combining empirical wavelet transformation (EWT) with ESN for financial forecasting was proposed by Gao et al. [15]. In that research, they tested the proposed approach with multiple datasets, including stock market indices, airline passengers, and other economic time series. They compared their approach with conventional ESN, Support Vector ESN, Laplacian ESN, wavelet denoising ESN, MLP, Support Vector Regression, and ARIMA models. The proposed EWT-ESN achieved superior performance across most datasets based on RMSE, MAPE, and MASE metrics [15].

Although significant machine learning approaches have been increasingly used in financial forecasting, there is still a research gap in the use of RC, specifically ESN, for exchange rate prediction. While ESNs have demonstrated promising results in stock market forecasting [14], [15], their full potential to model complex, non-linear dynamics of foreign exchange rates remains unexplored. In addition, there remains a significant research gap in exploring ensemble methods that specifically combine the strengths of LSTM, GRU, and RC. By utilizing the aforementioned approaches, this study tries to address these gaps by creating multiple ensemble models to test against one another.

3. Methodology

This section provides a detailed description of the methodology in this study. Fig. 1 presents a comprehensive visual overview of our methodological framework. This figure illustrates the workflow from data collection to the implementation of hybrid ensemble models, followed by the evaluation. Furthermore, a step-by-step explanation is provided for each process.

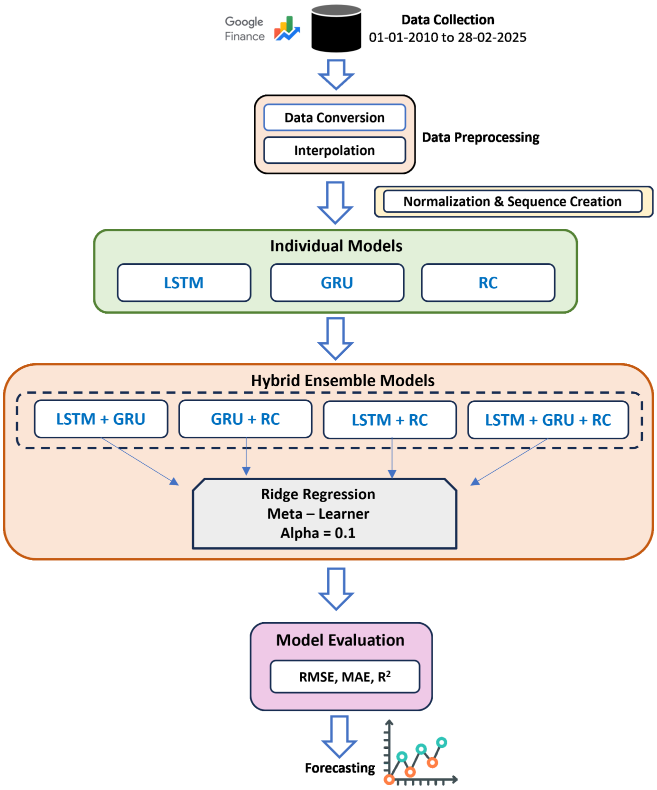

Figure 1. Exchange rate forecasting methodology

3.1. Data Collection: Retrieving USD to CAD Exchange Rates

Data are the fundamental component of any ML or DL process. For this study, USD to CAD exchange rate data was retrieved from Google Finance [17]. It is a financial platform provided by Google, which consists of real-time information related to stocks, indices, and currency exchange rates [17]. Furthermore, the data from 1st January 2010 to 28th February 2025 were accessed directly within Google Sheets using the GOOGLE FINANCE [18] function. Moreover, the dataset was saved in a comma-separated values (CSV) file. Only the date and close price attributes were preserved to facilitate further analysis.

3.2. Data Pre-Processing

To ensure accurate time-series analysis, formatting the date column is a significant task. During this stage, the date column is converted from its default format (object (string)) to DateTime format. Moreover, it was observed that the CSV file did not contain all daily closing values within the specified date range due to the closure of financial markets on weekends and public holidays. It is essential to have continuous values for getting good results in forecasting. To achieve this, missing dates were initially identified, and the respective exchange rates were placed using linear interpolation.

3.3. Data Exploration

With the intention of understanding and interpreting the dataset, several plots were generated during this phase. These visualizations helped in identifying trends and patterns within the close price of the USD/CAD exchange rate. Initially, a histogram (Fig. 2) was generated to visualize the distribution of close price values throughout the dataset. This histogram helps in identifying the most frequent exchange rate levels throughout the observed period.

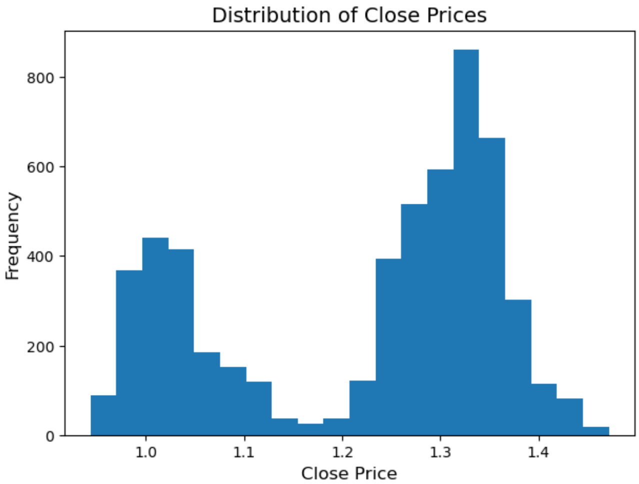

Figure 2. Distribution of close prices

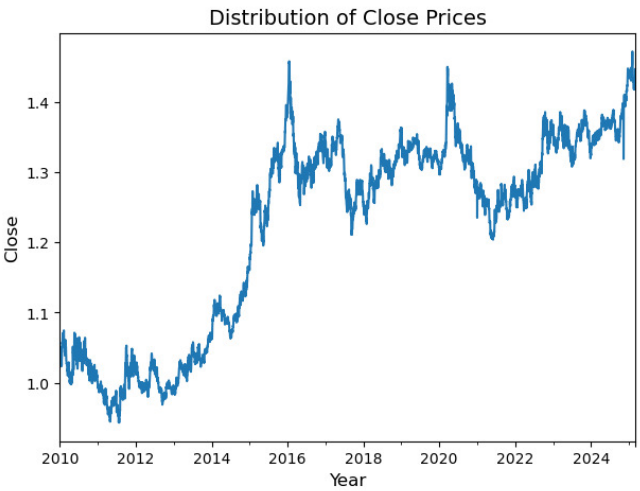

Figure 3. Daily closing price over time

The line plot of the daily closing price over time is depicted in Fig. 3. It allows us to observe long-term trends, overall price movement patterns, and volatility in the exchange rate. This plot also shows the daily closing price of the USD/CAD exchange rate over time, with the year on the x-axis and the close price on the y-axis.

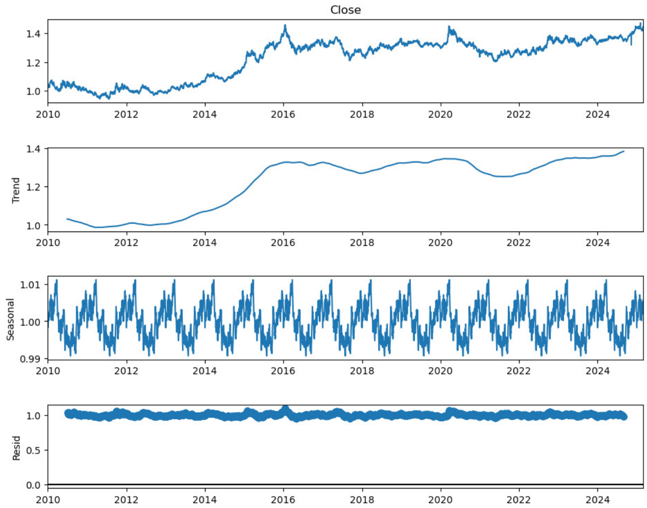

Figure 4. Seasonal decomposition of the close price time series

Moreover, to understand the underlying components of the USD/CAD exchange rate time series, a seasonal decomposition plot (Fig. 4) was obtained. This process separates the close price series into trend, seasonal, and residual components.

3.4. Data Preparation: Normalization and Creation of Sequential Data for Time Series Forecasting

During this process, the USD-CAD exchange rate data was first normalized using MinMaxScaler to bring all values within the range of 0 to 1. Then, the sliding window approach was utilized to transform the linear time series into a sequence-based format. In detail, continuous time series data were converted into overlapping input-output sequence pairs. This method creates input sequences of a predetermined length (30 days in our implementation), with each sequence serving as a feature set for predicting the subsequent time step.

3.5. LSTM, GRU, and RC Models: Architectures and Implementation

This subsection explores the three individual DL models implemented in this study. First, LSTM was implemented, which excels at capturing long-range dependencies through specialized memory cells. LSTM's standard architecture consists of three gates, namely forget, input, and output [12]. In detail, the forget gate serves as a filter, deciding on what information should be eliminated from the cell state. Moreover, it elects the new information to be preserved within the cell state. Also, the output gate determines what information should actually be produced based on what's in the cell state. As shown in Fig 5, these gates regulate information flow through the cell state. This mechanism enables the model to selectively remember or forget information across long sequences. [12].

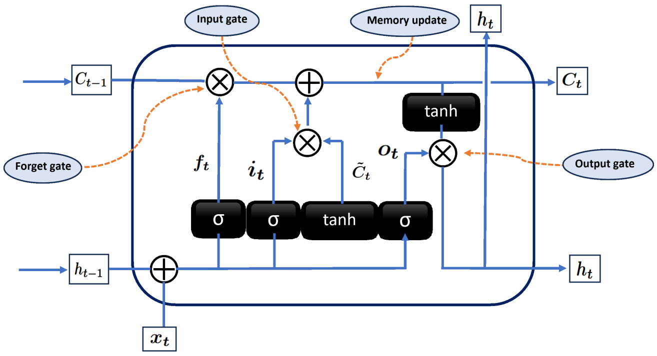

Figure 5. LSTM architecture with gates showing the information flow through cell states in which x_t represents the input at timestep t; h_t and h_{t-1} denote the current and previous hidden states respectively; C_t and C_{t-1} stand for the current and previous cell states; f_t indicates the forget gate output; i_t refers to the input gate output; o_t signifies the output gate output; $\tilde{C}_t$ represents the candidate cell state. For operations: σ denotes the sigmoid activation function; tanh refers to the hyperbolic tangent activation function; ⊕ indicates point-wise addition; ⊗ represents point-wise multiplication.

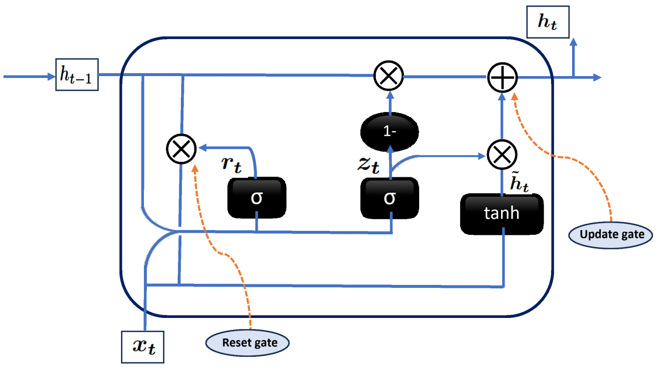

Figure 6. GRU architecture with gates showing the simplified recurrent structure in which x_t represents the input at timestep t; h_t and h_{t-1} denote the current and previous hidden states respectively; r_t indicates the reset gate output; z_t refers to the update gate output; $\tilde{h}_t$ represents the candidate hidden state. For operations: σ denotes the sigmoid activation function; tanh refers to the hyperbolic tangent activation function; ⊕ indicates point-wise addition; ⊗ represents point-wise multiplication; (1−) signifies the complementary value.

Following the LSTM's implementation, GRU was implemented. It offers a more streamlined architecture while maintaining strong performance on sequential data. As shown in Fig. 6, GRU contains only two gates: the reset gate and update gate [12]. The reset gate controls which parts of the previous information should be forgotten when creating new memory content. Meanwhile, the update gate acts like a switch that decides how much of the previous memory should be kept and how much new information should be added [12]. This streamlined design helps GRU to effectively process sequential data with fewer computational resources. Model configuration and training parameters used to configure the LSTM and GRU are mentioned in Table 1.

As the final individual model, RC was implemented. It is a distinct approach to recurrent neural networks that offers computational efficiency. As shown in Fig. 7, RC consists of three essential components: an input layer, a reservoir containing randomly connected neurons, and an output layer [14]. Unlike LSTM and GRU, the reservoir weights remain fixed after initialization, with only the output weights being trained. This distinctive characteristic significantly reduces computational complexity [13], [14], [15], during training while preserving the model's capacity to capture temporal patterns. The reservoir functions as a high-dimensional, nonlinear representation of the input where each neuron's state is determined by both the current input and the previous reservoir state [14]. Model configuration and training parameters used to configure the RC are mentioned in Table 2.

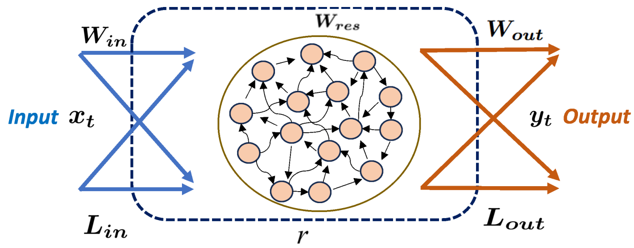

Figure 7. Reservoir Computing architecture showing the three main components. Key notations: x_t: input vector at time step t; y_t: output vector; r: reservoir with recurrent connections; W_{in}: input weight matrix connecting inputs to reservoir neurons; W_{res}: reservoir weight matrix (fixed random weights); W_{out}: output weight matrix (trained); L_{in}: input layer; L_{out}: output layer.

Moreover, after the implementation of the individual models, hybrid ensemble models were implemented. These hybrid approaches were facilitated to leverage the relative contributions of each base model through a meta-learning process.

3.6. Hybrid Ensemble Model Implementation

In most cases, rather than validating a prediction or forecasting using an individual approach, it is better to go for a hybrid approach. During this study, hybrid models combine predictions from two or more base models (LSTM, GRU, and RC) using a meta-learning framework. For each hybrid configuration, ridge regression (with $\alpha_{regularization} = 0.1$) was utilized

Table 1. Model configuration and training parameters for LSTM and GRU

Parameter	LSTM	GRU
Sequence length	30 days	30 days
Hidden units	50 units per layer	50 units per layer
Number of hidden layers	2	2
Activation function	ReLU	ReLU
Dropout	0.2	0.2
Optimizer	Adam	Adam
Learning rate	0.001	0.001
Loss function	MSE	MSE
Batch size	32	32
Epochs	50	50
Validation split	0.1	0.1
Train-test split	80%/20%	80%/20%

Table 2. Model configuration and training parameters of RC model

Parameter	RC
Sequence length	30 days
Reservoir neurons	200
Spectral radius	0.95
Sparsity	0.1
Noise	0.001
Activation function	tanh (reservoir)
Readout method	Ridge regression
Ridge alpha	1e-6
Loss function	MSE
Train-test split	80%/20%
Random state	42

as the meta-learner to optimally weight the predictions from the individual models. This approach allows the ensemble to benefit from the different temporal pattern recognition capabilities of each model type. Four hybrid combinations of models were implemented for further analysis and testing: LSTM+GRU, LSTM+RC, GRU+RC, and LSTM+GRU+RC.

The meta-learning approach works by first generating predictions from each individual model. These predictions are then used as features for the ridge regression model. From that, the model learns the optimal combination weights to minimize prediction error. This method represents a valuable alternative to simple averaging. It allows the meta-learner to assign different importance to each individual model based on its predictive performance, rather than treating all models equally.

3.7. Evaluation metrics: RMSE, MAE, and R^2

To evaluate the performance of our exchange rate prediction models, we employed several standard metrics that quantify prediction accuracy from different perspectives. These metrics allow for comprehensive model assessment and comparison.

Root Mean Squared Error (RMSE). RMSE is the square root of Mean Squared Error (MSE), providing an error metric in the same units as the original data:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{1}$$

where y_i represents the actual observed value, $\hat{y}_i$ denotes the predicted value, and n is the total number of observations. In this scenario, RMSE is more interpretable than MSE since it expresses the average model prediction error in the original units of the USD-CAD exchange rate. As with MSE, lower RMSE values indicate better prediction accuracy.

Mean Absolute Error (MAE). MAE quantifies the average magnitude of errors regardless of their direction:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{2}$$

where y_i represents the actual observed value, $\hat{y}_i$ denotes the predicted value, and n is the total number of observations. In contrast to RMSE, MAE assigns equal weight to all errors instead of placing greater emphasis on larger errors. This characteristic makes MAE more robust against outliers and offers a direct measurement of average error magnitude.

Coefficient of Determination (R^2). R^2, measures the proportion of variance in the dependent variable that is predictable from the independent variables:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{3}$$

where y_i represents the actual observed value, $\hat{y}_i$ denotes the predicted value, $\bar{y}$ is the mean of the observed values, and n is the total number of observations. R^2 ranges from 0 to 1, with higher values indicating that a larger proportion of the variance is explained by the model. R^2 of 1 indicates perfect prediction, while 0 suggests that the model does not explain any of the variability in the target variable.

4. Results and Discussion

This section presents the results obtained from the exploratory data analysis and model evaluation.

4.1. Results from Data Exploration

The exploratory analysis of the closing price revealed a distinctive bimodal distribution, as shown in Fig. 2. This histogram illustrates two distinct peaks: the first smaller peak is centered around 1.0-1.1, and the second larger peak is centered around 1.3-1.35. This distribution revealed that, in most cases, the exchange rate moves between these two price levels rather than relying on a single value. Moreover, Fig. 3 and the trend component (2nd panel) in Fig. 4 clearly show a significant jump in exchange rates occurring around 2014-2015. Also, since 2015, prices have mostly stayed in the higher range, with some ups and downs. From 2022 to 2024, a clear upward trend in the exchange rate was observed. The seasonal component (3rd panel) in Figure 4 shows consistent cyclical patterns throughout the entire period, indicating regular fluctuations that repeat annually. The residual component remains relatively stable around 1.0, suggesting that most price movements are well explained by the trend and seasonal patterns.

4.2. Model Performance Analysis

This section presents the performance analysis of individual models in predicting exchange rates, as well as the effectiveness of hybrid approaches. The results are supported by relevant plots and appropriate evaluation metrics mentioned earlier.

LSTM Model Performance Results. Fig. 8 illustrates the LSTM's notable predictive capability of forecasting exchange rates throughout the test data. In detail, the model follows certain actual exchange rate movements along with some areas where the prediction does not perfectly match the actual values. To further validate the aforementioned results, the higher RMSE value and higher MAE for LSTM shown in Table 3 indicate that this model produces larger average prediction errors.

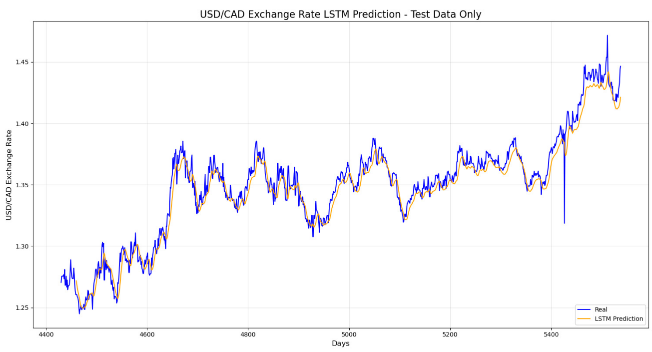

Figure 8. Prediction performance of LSTM on test data

GRU Model Performance Evaluation. According to the plot of the GRU in Fig. 9, it demonstrates higher predictive capability compared with LSTM. Moreover, GRU shows a better capacity in predicting currency fluctuations than LSTM. In addition to the observations, the performance metrics from Table 3 show lower error metrics (RMSE of 0.0063 and MAE of 0.0047) and a higher coefficient of determination (R^2 of 0.9762) for the GRU model. These improved metrics, when compared to the LSTM model, denote a significant enhancement in prediction accuracy achieved by the GRU approach.

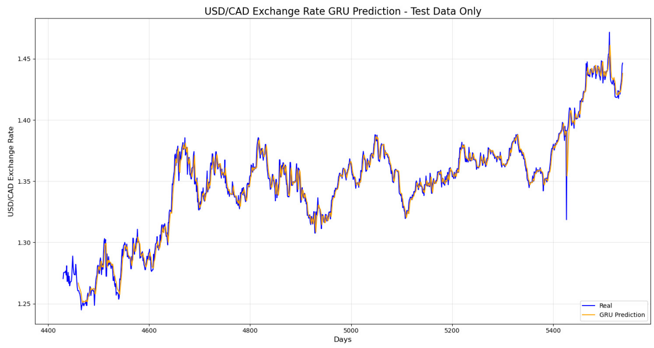

Figure 9. Prediction performance of GRU on test data

Reservoir Computing Evaluation. Fig. 10 illustrates the performance of the RC model in predicting USD/CAD exchange rates on test data. It struggles to align with actual trends and rapid fluctuations at early stages in predicting the exchange rate. Once the reservoir stabilizes, it appears to capture the underlying trends of currency exchange movements more effectively than the other recurrent architectures. Additionally, the results in Table 3 clearly demonstrates that RC provided exceptional performance metrics across all evaluation criteria. The RC model achieved the lowest RMSE and MAE values among all tested approaches. Furthermore, it displays the highest R^2 values, indicating superior predictive accuracy. These comprehensive results confirm that RC consistently outperformed individual models like LSTM and GRU. From the aforementioned results, it was noticed that

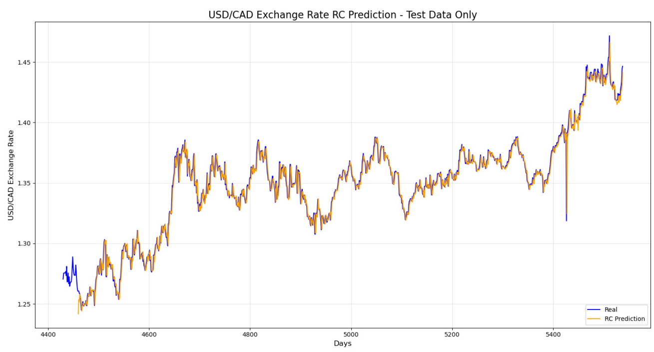

Figure 10. Prediction performance of RC on test data

there is room for further improvement in terms of reducing the RMSE and MAE values. With the intention of reducing the error values, hybrid approaches were utilized and tested against the same data. In the upcoming section, the results obtained from the hybrid approaches were discussed briefly.

Table 3. Performance comparison of individual and hybrid models on test data

Model	RMSE	MAE	R^2
LSTM	0.0091	0.0071	0.9515
GRU	0.0063	0.0044	0.9762
RC	0.0061	0.0038	0.9777
LSTM+GRU	0.0076	0.0057	0.9663
LSTM+RC	0.0060	0.0040	0.9786
GRU+RC	0.0061	0.0040	0.9777
LSTM+GRU+RC	0.0061	0.0040	0.9780

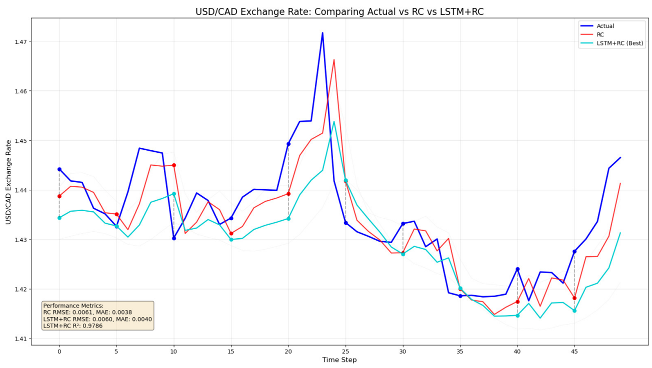

Figure 11. Prediction performance of LSTM+RC and RC on test data (last 50 time steps)

Hybrid Ensemble Model Effectiveness. During this analysis, individual models were combined and tested against one another in predicting the exchange rate using the close price. From Table 3, the hybrid model combining LSTM+RC yielded the best performance in terms of R^2 value (0.9786) and the lowest RMSE value (0.0060). However, a significant trade-off emerges when examining MAE: the RC model alone achieved the lowest MAE (0.0038), outperforming all hybrid combinations including LSTM+RC (MAE=0.0040). This suggests that although the hybrid approaches better capture the general trend, the RC model alone achieves the most accurate absolute predictions on average.

This performance trade-off is further illustrated in Fig. **??** over the final 50 time steps of the test period. Based on the figure, LSTM+RC achieved an RMSE of 0.0040, while RC had an RMSE of 0.0049. Therefore, LSTM+RC resulted in a lower total squared error. However, RC maintained a lower MAE (0.0040 vs. 0.0049 for LSTM+RC), indicating that the average difference from the actual values was smaller for RC. Both models tracked the actual U.S. dollar/Canadian dollar exchange rate very well. The results indicate that LSTM+RC is better at capturing variance and extremes (as measured by R^2 and RMSE), while RC is better at providing consistent average accuracy (as measured by MAE). The choice between these models depends on the application's needs: if minimizing large prediction errors is most important, LSTM+RC should be used; if maximizing consistency in average accuracy is the priority, RC is preferable. Because currency exchange rates are highly variable and both stability and accuracy are desirable, both models are good options, each with unique advantages.

5. Conclusion

Prediction of foreign exchange rates is a complex task due to the market's complex, volatile, and highly fluctuating nature. Researchers worldwide are continuously developing and discovering new forecasting models for this task. This study explored and adapted LSTM and RC's ESN models for forecasting the USD to CAD exchange rate. Experimental results revealed that individual models such as LSTM, GRU, and ESN performed reasonably well. However, the proposed hybrid model combining LSTM + RC outperformed them in terms of all the evaluation metrics considered. In addition, the hybrid approach highlights the benefits of integrating memory-based sequential learning with dynamic reservoir processing. Moreover, in the research, the proposed approach was tested only utilizing the closing price as a feature against the dates. Future research could further enhance prediction performance by incorporating additional influential factors and accommodating advanced techniques like attention mechanisms and transformer-based architectures.

References

[1] Z. Darvas and Z. Schepp. Exchange rates and fundamentals: Forecasting with long maturity forward rates. *Journal of International Money and Finance*, 143:103067, 2024. doi:10.1016/j.jimonfin.2024.103067. URL https://doi.org/10.1016/j.jimonfin.2024.103067. [Online; accessed 2024-05-01].

[2] S. Ranjit, S. Shrestha, S. Subedi, and S. Shakya. Comparison of algorithms in foreign exchange rate prediction. In *2018 IEEE 3rd International Conference on Computing, Communication and Security (ICCCS)*. IEEE, 2018. Accessed: 2025-05-01.

[3] Blueberry Markets. What Are the Factors Affecting USD/CAD?, 2024. URL https://blueberrymarkets.com/market-analysis/what-are-the-factors-affecting-usdcad/. [Online; accessed 2025-05-01].

[4] H. Makika, J.M.T. Romano, and R. Ballini. Deep learning techniques for time series forecasting: An application

for exchange rates. In *2023 IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2023. doi:10.1109/MLSP55844.2023.10285930. URL https://doi.org/10.1109/MLSP55844.2023.10285930. [Online; accessed 2025-05-05].

[5] M.S. Islam and E. Hossain. Foreign exchange currency rate prediction using a gru-lstm hybrid network. *Soft Computing Letters*, 3:100009, 2021. doi:10.1016/j.socl.2020.100009. Accessed: 2025-05-01.

[6] M. Qi and Y. Wu. Nonlinear prediction of exchange rates with monetary fundamentals. *Journal of Empirical Finance*, 10:623–640, 2003. doi:10.1016/S0927-5398(03)00008-2. URL https://www.elsevier.com/locate/econbase.

[7] C. Cheng, A. Sa-Ngasoongsong, O. Beyca, T. Le, H. Yang, Z. Kong, and S.T.S. Bukkapatnam. Time series forecasting for nonlinear and non-stationary processes: A review and comparative study. *IIE Transactions*, 47(10):1053–1071, 2015. doi:10.1080/0740817X.2014.999180. URL https://doi.org/10.1080/0740817X.2014.999180. [Online; accessed 2024-05-01].

[8] E. Colombo and M. Pelagatti. Statistical learning and exchange rate forecasting. *International Journal of Forecasting*, 36(4):1260–1289, 2020. doi:10.1016/j.ijforecast.2019.12.007. URL https://doi.org/10.1016/j.ijforecast.2019.12.007. [Online; accessed 2024-05-05].

[9] Bank of Canada. In Focus 2: Recent Factors Affecting the Canada–US Exchange Rate, 2025. URL https://www.bankofcanada.ca/publications/mpr/mpr-2025-01-29/in-focus-2/. [Online; accessed 2025-05-01].

[10] Investopedia. Loonie: Meaning, Overview, Role in Economy, 2022. URL https://www.investopedia.com/terms/l/loonie.asp. [Online; accessed 2025-05-01].

[11] R. Dev, M.N. Alam, and M.L. Rahman. Arima model based time series modelling and prediction of foreign exchange rate against us dollar. *International Journal of Recent Technology and Engineering*, 11(3):100–106, 2022. doi:10.35940/ijrte.C5792.111322. URL https://doi.org/10.35940/ijrte.C5792.111322. [Online; accessed 2025-05-05].

[12] A. Lawi, H. Mesra, and S. Amir. Implementation of long short-term memory and gated recurrent units on grouped time-series data to predict stock prices accurately. *Journal of Big Data*, 9, 07 2022. doi:10.1186/s40537-022-00597-0.

[13] H. Jaeger. The "echo state" approach to analysing and training recurrent neural networks. GMD Report 148, German National Research Center for Information Technology, 2001.

[14] W.-J. Wang, Y. Tang, J. Xiong, and Y.-C. Zhang. Stock market index prediction based on reservoir computing models. *Expert Systems with Applications*, 178:115022, 2021. ISSN 0957-4174. doi:10.1016/j.eswa.2021.115022. URL https://www.sciencedirect.com/science/article/pii/S0957417421004632.

[15] R. Gao, L. Du, O. Duru, and K.F. Yuen. Time series forecasting based on echo state network and empirical wavelet transformation. *Applied Soft Computing*, 102:107111, 2021. doi:10.1016/j.asoc.2021.107111. URL https://doi.org/10.1016/j.asoc.2021.107111. Accessed: 2025-05-05.

[16] L. Barua and M.J. Uddin. Exchange rate forecasting via generative adversarial network (gan). In *2024 IEEE International Conference on Computing, Applications and Systems (COMPAS)*. IEEE, 2024. doi:10.1109/COMPAS60761.2024.10796492. URL https://ieeexplore.ieee.org/document/10796492. [Online; accessed 2025-05-05].

[17] Google Finance. Google Finance, 2025. URL https://www.google.com/finance/. [Online; accessed 2025-03-14].

[18] Google Support. GOOGLEFINANCE function, 2025. URL https://support.google.com/docs/answer/3093281?hl=en. [Online; accessed 2025-03-14].