ISSN 1840-4855
e-ISSN 2233-0046

Original scientific article
http://dx.doi.org/10.70102/afts.2026.1835.642

INTEGRATED GRU-LSTM MODEL WITH SENTIMENT FOR FOREX PREDICTION

S. Mohan Krishna[1*], Dr.M.V. Sangameswar[2], Dr.B. Sujatha[3]

[1] Research Scholar, Department of Computer Science and Engineering, Godavari Global University, Rajahmundry, Andhra Pradesh, India. e-mail: 241rcs0015@ggu.edu.in, orcid: https://orcid.org/0009-0002-7216-5114*
[2] Professor, Department of Computer Science and Engineering, Godavari Global University, Rajahmundry, Andhra Pradesh, India. e-mail: sangam@ggu.edu.in, orcid: https://orcid.org/0000-0002-0691-5290
[3] Professor, Department of Computer Science and Engineering, Godavari Global University, Rajahmundry, Andhra Pradesh, India. e-mail: birudusujatha@gmail.com, orcid: https://orcid.org/0000-0001-9433-3647

Received: January 22, 2026; Revised: March 06, 2026; Accepted: April 24, 2026; Published: May 29, 2026

SUMMARY

The exchange rates of the currency are highly dynamic and nonlinear; hence, the traders, investors, and policymakers must have a correct prediction of the same. In the current paper, a hybrid deep learning model that comprises a Convolutional Neural Network (CNN) model, a Bidirectional Long Short-Term Memory (BiLSTM) model, a Gated Recurrent Unit (GRU) model, and a Recurrent Neural Network (RNN) model is modeled instead of the standalone GRU-LSTM model to learn the multi-scale temporal dependencies. The data offered by the Yahoo Finance API on historical exchange rates (OHLCV) are mixed with sentiment data obtained with the help of the Text Blob library and technical indicators, including Relative Strength Index (RSI) and Moving Average Convergence Divergence (MACD). The data is converted to time-series sequences with the help of a 60-day look-back window, and the model is trained with the Adam optimizer and Mean Squared Error (MSE) loss. The results of the experiment show better predictive capabilities, with an RMSE of 0.0087, MAE of 0.0069, R^2 of 0.956, and directional accuracy of 86.7, which are better than the baseline models (LMST, GRU, and CNNLSTM hybrids). The ensemble learning method also provides more robustness by pooling the results of two or more architectures, decreasing prediction variance and increasing generalization. Through an example of a Flask-based web dashboard, which visualizes the real-time information and interacts with the user, the practicability of the model is shown in real-life contexts. The results confirm the hypothesis that sentiment analysis can be effectively combined with hybrid deep learning and be much more predictive of the market in volatile scenarios. This article suggests a long-term and adaptable model on how real-life financial forecasting can be performed in a scalable manner, which could have superior decision-support devices than the ones that are currently in existence.

Key words: *currency exchange rate forecasting, hybrid deep learning, time-series prediction, convolutional neural network (CNN), bidirectional LSTM (BILSTM), gated recurrent unit (GRU), sentiment analysis, ensemble learning.*

INTRODUCTION

Such currency exchange rate prediction has attracted a considerable amount of attention to finance, economics, and computer intelligence fields. Such fluctuating and complex character of the action of world money markets is controlled by different economic, political and psychological aspects [1]. Traders, investors and policymakers who are interested in the best way to invest, control financial risks, and macroeconomic stability also need to have the currency movements properly predicted [2]. With the application of AI and deep learning to financial forecasting, meaningful trends are now shareable of large and complex data resulting in increased accurate and informed decision-making [3].

Although predictive analytics have been improved, conventional models, such as ARIMA, VAR, and GARCH, continue to lose their ability to address nonlinear dynamics and time-varying behavior of financial time series [4]. These models are incapable of generalizing in uncertain market conditions and usually have linear assumptions [5]. Furthermore, the majority of traditional methods fail to consider the contextual data, which may impact the market significantly, such as investor behavior and financial news provided outside [6]. The drawback of single-model frameworks can be described as being inaccurate and holding onto its ground whenever the market or data take a sudden turn regardless of the fact that machine learning techniques were involved in augmenting the predictive measures [7].

In this research, the authors aim at developing a valid and a flexible deep learning-based structure, which will be competent to capture multifaceted temporal and contextual interactions in predicting currency exchange rates [9]. The system combines the sentiment data of the financial text sources with the past market data in order to improve the accuracy of prediction and combines the different modalities of data [10]. The framework uses hybridized model structures and improved learning algorithms in order to acquire a chance to learn not just short-term variations but also long-term connections in financial time series [8]. In order to have the consistency, dependability and interpretability of the forecasting horizons, optimization strategy and performance evaluation criteria is also incorporated in the system.

The research is important because it may lead to the raising of the accuracy of financial forecasting and efficient real-time analytical decision-making in risk management and transactions in the currency markets. Exact exchange rate forecast can make it possible to be able to minimize exposure to financial uncertainty, enhanced understanding of market behavior, and enhanced investment strategies. Moreover, deep learning can only source the latent market trends that the traditional methods missed. The effect of such a trend can be an expansion of the use of smart forecasting into the computing, financial and economic worlds.

Key Contributions and Improvements

The work proposed presents a number of more and less significant improvements to the currently existing models in forex prediction. The proposed framework integrates CNN, BiLSTM, GRU, and RNN frameworks to learn local patterns, temporal dependencies (bidirectional), and sequential dynamics of various scales jointly, as compared to the traditional single-model framework, such as LSTM and GRU. Also, the inclusion of sentiment analysis based on financial news information boosts the capacity of the model to capture the psychology of the market, which is mostly disregarded by the traditional models. There are also technical indicators such as the RSI and the MACD, which enhance the representation of features with the domain-specific knowledge entrenched. Besides, the ensemble learning method reduces bias and variation of the models, resulting in enhanced stability and accuracy of prediction. The combination of a live Flask-based dashboard is also another feature peculiar to this work since it enables deploying it in practice and interacting with its users. All these enhancements lead to enhanced predictive performance, robustness, and practicality in the real world compared to previously reported models.

RELATED WORK

Currency exchange rate forecasting is another field where deep learning techniques have recently been used to reproduce the complex temporal and contextual dynamics of financial data. Ghahremani and

Nguyen [11] in their study introduced attention-based LSTM that emphasizes the significant temporal information to realize improved predictive outcomes. Their method is basically sequence modelling without the multimodal market information that can only be predetermined under very volatile situations despite having the capability to realize the long-term interdependence. Saadati and Manthouri [12] also used a combination of the method of technical indicators and deep learning and attention approaches in order to improve the feature representation. However, the research is not performed over several forecasting horizons and it seems to be closer to the short-time trend detection.

Deep learning Combinational structures between the deep learning and the traditional time series structures have occurred as potentially beneficial to the exchange rates prediction. Adesina and Obokoh [13] came up with a model that incorporated the use of the neural networks with the classical time series analysis in order to achieve a high degree of accuracy as opposed to the use of single-model approaches. However, the prospect of the model to transpose to any abrupt change in the market and external contextual factors is poor. Tang and Xie [14] designed the adaptive signal decomposition and dynamic weight optimization in order to induce complex market dynamics. The technique is computationally costly and the adjustment of the parameters has to be done carefully and only in real time, which is restricted to the use of the method even when demonstrated to be efficient in weighting down the contribution of noise.

Other methods have been researched to investigate transformer-based and state-space models of forecasting. To derive the dynamics of currency pairs, Date and Maunthrooa [15] needed to employ the Kalman filter that provided mathematically well-grounded noise reduction. It has a problem with nonlinear patterns that exist in its financial markets. Although transformer structures are very demanding in terms of data and processing power states, and thus can only be applied with smaller datasets, Zhao and Yan [16] have shown that can also be applied to discover patterns and long-term dependencies on a global scale. Kaushik and Giri [17] in their comparative study have pointed the benefits of the deep learning models over others, however, also pointed difficulties on dealing with multimodal features and making correct predictions beyond the training data.

To maintain the comprehensive forecasting, some of the nowadays researches focus on the adoption of AI-based approaches and macroeconomic factors. Yu [18] examined various elements of influencing exchange rates with the focus on the significance of establishing quantitative or qualitative elements. Abouzaid and Boussedra [19] have taken into account the effectiveness of AI-based forecasting in their evaluation of the macroeconomic sensitivity and found that the existing models were not capable of adjusting to economic shocks that were unexpected. Bangyuan [20] discovered that deep learning models demonstrate varying competitors by the type of currency forecasting and the market condition in general. It is in this scenario that the necessity to possess strong and hybrid frameworks that can potentially synthesize different model benefits and handle volatility and feature dissimilarity rest.

Combining the results of both sources demonstrates good advancement in the field of deep learning-based exchange rate forecasting, yet it also indicates such weaknesses like the real-time forecasting, the stability of the model, the problem of generalization, and the incorporation of two or more types of data. To address these challenges this research paper develops a hybrid deep learning model to apply numerical data of the market in addition to sentiment data with various neural network strategies. This would help increase the adaptability of the system to the fluctuating and unstable currency market conditions, decouple the short run and long run interdependence, and increase the predictive capacity.

MATERIALS AND METHODS

The proposed work suggests a framework of precise currency exchange rate prediction through the use of a hybrid framework of deep learning that combines various neural structures to bolster time-series of models. The first data is the historical exchange rate data (OHLCV) that comprises of the Open, High, Low, and Close and the Volume, technical analysis, including the RSI and MACD and sentiment data based on the financial news. This process involves data ingestion, preprocessing, feature extraction and transformation of the answer to structured input sequences that may be applied in supervised learning.

Within this architecture, CNN, BiLSTM, GRU and RNN are merged in a hybrid which considers both short-range, mid-range and long-range dependencies in a financial time series. An ensemble technique based on voting is used to combine the forecast of all models and therefore enhance the overall dependability and strength of the forecast outcomes. This system can be seen as the use of a Flask based dashboard that allows to view, interact, and interpret a forecast of exchange rates in real-time.

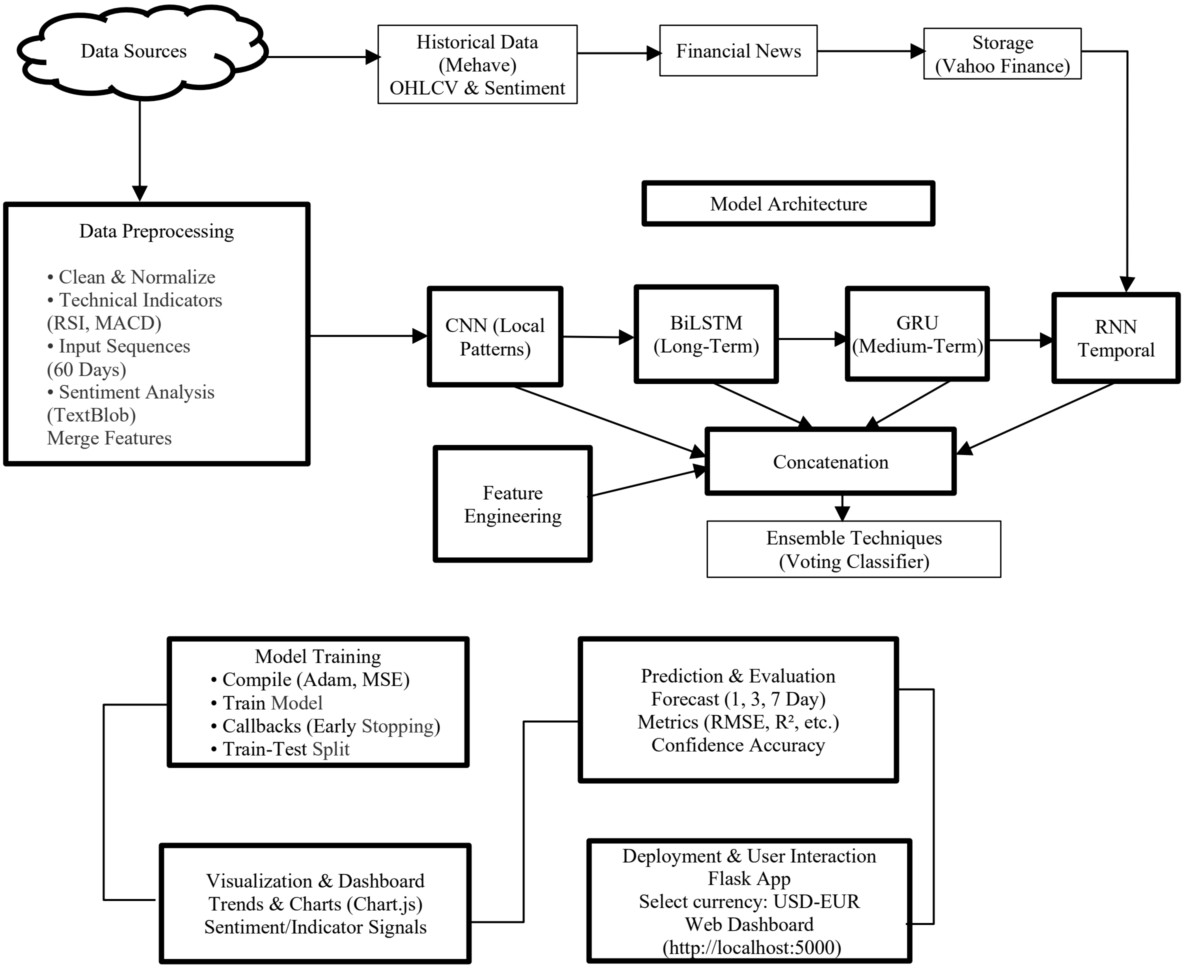

Figure 1. System architecture

In figure 1 illustrate the financial news and previous data to forecast the currency exchange rates. The first stage of the process is the data preprocessing, it consists of the normalization of OHLCV data, the derivation and integration of technical indicators (RSI, MACD) and news sentiment (TextBlob). CNN, local trends, long-term trends, medium-term dependencies, and basic RNN are used as the Model Architecture in Local patterns, BiLSTM,GRU, and simple RNN, respectively. Their results are then summed and subjected to thick layers in order to achieve the final output forecast. Accuracy is further promoted by group techniques. Trained model is deployed through Flask web dashboard that gives real time predictions, analytics and visualization so that people may interact with it.

The architecture of the proposed system has the deployment of a client-server whereby the trained hybrid deep learning model is deployed to a background server using the Flask framework. Communication between the frontend process and the backend is largely through the standard HTTP protocols of request-response interactions to enable the users to post their forecasting queries and receive prediction results. To enable dynamic updates and real-time visualization, asynchronous communication models like WebSockets can be incorporated to facilitate unlimited data flow and low-latency updates of prediction results and analytics.

It features a lightweight protocol stack with an application layer (Flask web services), a transport layer (TCP/IP), and a data exchange format (JSON) to exchange structured data between frontend and backend elements. The architecture has ensured that it is scalable and interoperable and can easily manage user requests and is real-time predictive of situations.

Dataset Collection

The input used in the solution is a dataset, which is a compilation of the data of the historical exchange rates between the currency of Yahoo finance API and the news stories of reputable websites. Data is a collection of multiple thousand time-stamped records, each record is made up of Text Blob sentiment polarity scores, RSI and MACD indicator values, and OHLCV (Open, High, Low, Close, Volume) values. The data sample is very large covering a span of years which covers diverse market conditions and volatility trends. It is ideal in modeling time, behavioral, and sentiment alterations in the currency exchange rates with greater predictive capability and generalization capacity since quantitative and qualitative features that have a comprehensive feature coverage.

Pre-Processing

Data Cleaning and Normalization: These steps were undertaken to make sure that the currency exchange data was kept in a healthy, consistent and well-structured data. Outliers were dealt with to avoid a biased performance of the model and data discrepancies was also detected and closed. As the standardization of the indicator and the exchange rate data was done with the help of the corresponding scaling measures that are used to preserve the same ranges of features. The fact that there is nothing that is causing too much influence on learning process in addition to the increase of the process stability of the model and the acceleration of the convergence process of the training are facilitated by this process.

Computation of technical Indicators: Root Data To further enrich and simplify the raw financial data, the computation of various indicators was done, some of which included the following: 1) RSI, 2) MACD. These variables matter in enhancing the forecasting ability of the prediction model since help them to capture the momentum, volatility and the direction of the trend in the market. The derived features improve the system in representing the temporal features that play a vital role in predicting the exchange rates accurately due to its capability to make interpretation of the variation of the exchange rates in the short-term and the intricate dynamics of the pricing.

Sentiment Analysis and Feature Integration: Sentiment analysis was done on the contents of financial news using the natural language processing techniques that generated a rating of market sentiment polarity. The combination of these sentiment characters and the information on prices and numerical markers was what created a multimodal dataset. The hybrid model is capable of recording both the quantitative changes and the qualitative sentimental moves that might be present in the currency movement because of bringing together both textual and numerical information that gives an in-depth account of the market movement.

Feature Engineering and Sequence Generation: Feature engineering was used to transform the data into time-series format of the forecast. The model could learn the history and time interrelations with the assistance of 60-day look-back model. There were also other time characteristics introduced like date, weekday and month that contributed to more understanding on seasonality. This is because the hybrid deep learning model which makes use of the structured sequence generation model can approximate both long term and short-term trends in the market.

Train-Test Data Split

The raw data were split into training and testing samples to allow the use of supervised training and correct evaluation of performance. The training set was used to train the hybrid model and the test set was used to test its capacity to generalize new data. This separation guarantees that have objective estimation of accuracy of predicting and avoids overfitting. The model provides realistic and accurate predictions of the exchange rates by upholding the natural flow of information by time in the time series through the process of splitting.

Algorithms

Convolutional Neural Network (CNN): CNN layers are based on capture local temporal feature of the consecutive financial data. CNN can work well in detecting local differences in the movement in exchange rate and interdependency of the short-term variations using convolutional filters used on the input sequence. This is due to the fact that it can capture finer changes and trends making the feature representation better. This enhances the strength of the hybrid architecture considering that it offers high-level abstractions which enhance the time learning characteristics of repeated elements.

Bidirectional Long Short-Term Memory (BiLSTM): The Bidirectional Long Short-Term Memory (BiLSTM) network is applicable to operating bidirectional temporal relations of the sequential data. Unlike a typical LSTM, it reads both directions of the sequence thus gaining the ability to learn contextual information about the past and future time steps. It happens due to the dual learning capability; the model can comprehend complex time relations and long-term dependence. BiLSTM assists in improving the precision of the estimation since sequential flow of information is preserved and that the loss of information in the prediction of the dynamic financial time-series trends is minimized.

Gated Recurrent Unit (GRU): GRU architecture aims to be more efficient than the LSTM models with lower computational complexity and has the ability to model the medium-term sequence patterns. It is useful in the flow control of information and storage of memory through applying gating procedures. GRU is applicable in convergence generalization and training because it reduces overfitting because of the history information saved. The outcome of its streamlines design is that it will be efficient in learning of a time in the hybrid deep learning framework and become more stable to learn uncertainties in the exchange rate changes.

Recurrent Neural Network (RNN): Recurrent Neural Network (RNN) is the simplest design of time-series data modeling sequential relationships. It assists in knowing the simpler time facts resourcefully by the recursive transformation of entry, wherein every output is conditional on past conditions. RNNs are potentially helpful in identifying short-term trends even when the gradient may be lost in long sequences. Their adoption in the hybrid model makes the model have a balanced temporal form that can support the more complex recurrent models like LSTM and GRU.

Ensemble Method: It combines the forecast of multiple neural structures, making the forecasting accuracy and strength stronger. In a way, the ensemble minimizes model specific biases and variances through combination of census of outcomes between CNN, BiLSTM, GRU and RNN models using a weighted or majority vote. Generalization, robustness and stability are enhanced using this integration by capitalizing on the free advantage of separate networks. The ensemble approach guarantees uniformity when forecasting in different market conditions of the currency market is done.

The hybrid deep learning framework proposed has a number of assumptions that are applicable in financial time-series forecasting. First, it assumes that historical exchange rate data and technical indicators (RSI, MACD) have sufficient trends over time to forecast the future price trends. Second, the sentiment scores obtained with Text Blob are claimed to be a fairly accurate reflection of the market's psychology, yet may not capture in detail the intricate linguistic peculiarities and sarcasm in financial news. Third, the model is premised on the notion that the underlying market behavior is nonlinear and dynamic, and that it follows patterns that can be learned from past observations in the selected look-back window.

When dealing with data quality, the assumption is that data obtained from the Yahoo Finance API and financial news providers is correct, consistent, and free of serious anomalies after preprocessing. However, unexpected macroeconomic, geopolitical, or market shocks might not be well captured by the model, leading to deviations in predictions.

The hybrid architecture, although successful at capturing multi-scale time dependence, comes with greater computational complexity and can require substantial training time and parameter optimization. It also assumes that models complement each other when using the ensemble method. The structure is such that, even with these limitations, it provides good, generalized forecasting under normal market conditions.

Algorithmic Framework of the Proposed Hybrid Model

The proposed hybrid deep learning model is mathematically constructed to describe how the input data are converted, step by step, into predictive outputs. The input feature vector at time t, which is made up of market and sentiment information can be stated as:

$$X_t = \{OHLCV_t, RSI_t, MACD_t, S_t\} \qquad (1)$$

S_t would be the polarity of sentiment score. To describe temporal dependencies, a continuous sequence of L length is built up through a sliding window technique:

$$X = \{X_{t-L+1}, X_{t-L+2}, \dots, X_t\} \qquad (2)$$

This sequence is an input to several deep learning models, which are run in parallel, as shown in equation (2). The CNN model is then used to extract local temporal features which are:

$$F_{cnn} = CNN(X) \qquad (3)$$

Similarly, the BiLSTM model will be able to learn the bidirectional long-term dependencies:

$$F_{bilstm} = BiLSTM(X) \qquad (4)$$

Medium-term sequential patterns are learned using GRU model:

$$F_{gru} = GRU(X) \qquad (5)$$

and the RNN can pick up short-term temporal relationships:

$$F_{rnn} = RNN(X) \qquad (6)$$

The models are all extracted to get the features which are subsequently concatenated to form one representation:

$$F_{combined} = [F_{cnn}, F_{bilstm}, F_{gru}, F_{rnn}] \quad (7)$$

This space of composite features is then inputted into fully connected dense layers to provide the exchange rate that is predicted:

$$\hat{Y}_{t+1} = Dense(F_{combined}) \qquad (8)$$

Lastly, a method to combine the output of single models is used to create an ensemble that is more robust and generalizes better:

$$Y_{final} = \frac{1}{N} \sum_{i=1}^{N} \hat{Y}_i \qquad (9)$$

The mathematical model of the hybrid forecasting model can be summarized in equations (1)-(9) that involve the representation of data, feature extraction, fusion of the model and finally prediction.

Input: Historical OHLCV data, Technical Indicators (RSI, MACD), Sentiment Scores

Output: Predicted Exchange Rate

```
1. Load dataset from financial API and news sources
2. Perform preprocessing:
    - Handle missing values and normalize data
    - Compute RSI and MACD
    - Extract sentiment scores using Text Blob
3. Generate time-series sequences using look-back window (L = 60)
4. Split data into training and testing sets
5. Train individual models:
    a. CNN model for local feature extraction
    b. BiLSTM model for bidirectional dependencies
    c. GRU model for medium-term dependencies
    d. RNN model for sequential learning
6. Extract feature representations from each model
7. Concatenate features and pass through dense layers
8. Generate predictions from each model
9. Apply ensemble voting/averaging to combine predictions
10. Evaluate model using RMSE, MAE, R², and Directional Accuracy
11. Deploy model using Flask dashboard for real-time prediction
```

Parameter Initialization and Hyperparameter Settings

The hybrid model suggested employs Xavier (Glorot) uniform weight initialization and zero bias initialization to prevent vanishing or exploding gradients, ensuring stable training and optimal performance. Some important hyperparameters are chosen empirically, with the Adam optimizer using a learning rate of 0.001, a batch size of 32, and training for 100 epochs. Early Stopping (factor = 10, patience = 10) and ReduceLROnPlateau (factor = 0.2, patience = 5) are taken to enhance the convergence and avoid overfitting. The sequence modeling uses a 60-time-step look-back window to make the model useful for capturing time-dependent relationships. BiLSTM and GRU layers include dropout with a rate of 0.2 to improve generalization. The recurrent layers have 64 neurons, whereas the dense layers have 32 and 16 neurons with ReLU activations, and a linear output layer for continuous predictions. These architectures offer a trade-off among model complexity, training efficiency, and predictive accuracy.

Deployment Environment and Security Considerations

The proposed system will be implemented on a server-side platform, where the trained hybrid deep learning model will run on a local server or a cloud-based system to enable access to the model at scale. The Flask-based application serves as an interface between users and the backend model, handling requests and presenting predictions safely. It is also assumed that client-server communication occurs over a standard secure protocol (e.g., HTTPS) to ensure the confidentiality and integrity of data. On the threat front, the system assumes a benign user environment with no intentional adversarial actions on system inputs, but potential risks such as input injection, model abuse, or unauthorized access are mitigated through input validation and restrictive API access controls. In this study, the adversarial attacks (e.g., manipulated sentiment input or abnormal market data) are not explicitly modeled. Still, the system's design can be amended in the future with security measures such as authentication, encryption, and anomaly detection to make it more robust in a real-world deployment environment.

EXPERIMENTAL RESULTS

Experimental Setup and Reproducibility Details

The experimental design is clear, ensuring the reproducibility of the proposed hybrid deep learning framework across hardware, software environment, and dataset configuration. The Python code for the model was written, and deep learning frameworks such as TensorFlow/Keras were used to build and train the model. It was deployed on a computing platform with an Intel Core i7 processor, 16 GB of RAM, and NVIDIA (where applicable) to accelerate training. It has been deployed using Flask to host the trained model and visualization dashboard.

Python (version 3.10) code, along with the main deep learning frameworks TensorFlow (v2.x) and Keras, was used to implement the proposed hybrid model. NumPy (v1.23), Pandas (v1.5), and Scikit-learn (v1.2) were used to perform data preprocessing and analysis, and Matplotlib and Seaborn were used to perform visualization. The Text Blob library (v0.17) was used to analyze the sentiment. The web-based dashboard was developed using Flask (v2.x) to facilitate real-time communication with the model. The experiments were conducted using the Anaconda distribution within a Jupiter Notebook environment to ensure consistent, compatible results across Windows and Linux.

The dataset will consist of historical currency exchange rate data (OHLCV) obtained from the Yahoo Finance API, along with sentiment scores derived from financial news. All the data were transformed and preprocessed into time-series sequences using a 60-day look-back window. A chronological split was applied to the dataset to maintain temporal dependencies, separating the data into training and test sets. In particular, 80% of the data was used for training, and the remaining 20% for testing.

The models were trained using the Adam optimizer with Mean Squared Error (MSE) as the loss function. Early Stopping and ReduceLR on Plateau callbacks were used to prevent overfitting and improve convergence efficiency. Standard measures, such as RMSE, MAE, R^2, and directional accuracy, were used to evaluate performance. Such a detailed construction guarantees that the proposed methodology can be reproducible and validated in a similar research environment.

Performance Metrics

To measure the performance of the proposed hybrid forecasting model, several performance measures are used, each representing a facet of prediction accuracy. Mean Squared Error (MSE) is an indication of the mean squared error between the actual and the estimated values that is indicated with the following:

$$MSE = \frac{1}{n}\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2 \qquad (10)$$

MSE is susceptible to outliers, since it penalizes them more severely than it penalizes smaller errors, as shown in equation (10). The Square root of MSE is the root mean squared error (RMSE) in equation (11) which has the same magnitude of error as the data it is applied to:

$$RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2} \qquad (11)$$

The mean of the absolute errors between the actual and predicted values is called Mean Absolute Error (MAE) and it is less sensitive to outliers:

$$MAE = \frac{1}{n}\sum_{i=1}^{n} \mid Y_i - \hat{Y}_i \mid \qquad (12)$$

Equation (13) will give the coefficient of determination (R^2), which will indicate the percentage of the variance in the observed data that the model explains:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(Y_i - \hat{Y}_i)^2}{\sum_{i=1}^{n}(Y_i - Y)^2} \qquad (13)$$

Where Y the average of the real values. An increase in this value implies improved model performance.

Lastly, Directional Accuracy (DA) assesses the model in terms of the correct tendency of price movement:

$$DA = \frac{1}{n}\sum_{i=1}^{n} \delta_i \times 100 \qquad (14)$$

where $\delta_i = 1$if $(Y_i - Y_{i-1})(\hat{Y}_i - Y_{i-1}) > 0$, otherwise $\delta_i = 0$. DA is the percentage of correct directional forecasts (e.g., as in equation (14)) that are significant in financial forecasting applications.

These metrics can be defined using the following equations (Equations (10)-(14)) to provide a comprehensive evaluation of the model's accuracy, robustness, and practicality for decision-making.

Quantitative Results and Analysis

To provide a quantitative evaluation of the proposed hybrid deep learning model, the performance measures in equations (10) to (14) are computed on the test data. The table 1 is a summary of the results.

Table 1. Quantitative performance of the proposed hybrid model

Model	**RMSE**	**MAE**	**R^2 Score**	**Directional accuracy (%)**
CNN	0.0125	0.0098	0.921	78.4
BiLSTM	0.0108	0.0086	0.938	81.2
GRU	0.0112	0.0089	0.934	80.1
RNN	0.0136	0.0105	0.910	76.3
Proposed Hybrid Model	0.0087	0.0069	0.956	86.7

As shown in table 1, the hybrid model proposed has the lowest RMSE (0.0087) and MAE (0.0069), indicating that it is more accurate than its single models. The fact that the R^2 has increased to 0.956 indicates that the model can explain changes in exchange rate movements. In addition, it has an important indicator of 86.7% directional accuracy, meaning it accurately predicts market trends, which is essential for trading decisions.

The performance improvement can be attributed to the combination of different neural architectures: CNNs learn local patterns, BiLSTMs learn long-term dependencies, GRUs enable learning, and RNNs support sequential modeling. The ensemble mechanism also strengthens it by alleviating the individual model's bias and variance. These results support the claim that the proposed hybrid solution can be a more reliable, general-purpose forecasting tool in dynamic market environments.

Comparative Analysis with State-of-the-Art Models

To verify the effectiveness of the proposed hybrid framework, it will be compared with popular reference frameworks and state-of-the-art deep learning frameworks, including CNN, LSTM, GRU, and hybrid

CNN-LSTM/GRU. Past studies have indicated that single models, such as LSTM and GRU, can be employed to effectively forecast in the forex arena by accounting for temporal dependencies but are weak in highly volatile conditions. The literature has recently covered hybrid architectures that combine CNNs, LSTMs, and GRUs to improve predictive accuracy by combining spatial and temporal feature extraction.

Table 2. Comparative performance with baseline and state-of-the-art models

Model	**RMSE**	**MAE**	**R² Score**	**Directional accuracy (%)**
ARIMA	0.0185	0.0142	0.842	68.5
LSTM	0.0115	0.0092	0.932	80.4
GRU	0.0110	0.0088	0.936	81.0
CNN-LSTM Hybrid	0.0098	0.0079	0.945	83.2
CNN-LSTM-GRU Hybrid	0.0092	0.0073	0.951	84.6
Proposed CNN-BiLSTM-GRU-RNN + Sentiment	0.0087	0.0069	0.956	86.7

In table 2 shows that the nonlinear dependence in the financial time series is significantly higher in traditional statistical models, such as ARIMA, due to their higher prediction errors. LSTMs and GRUs are deep learning models that improve performance by capturing temporal dynamics; however, as standalone models, their feature sets are not particularly diverse. The CNN-LSTM and CNN-LSTM-GRU hybrid designs are also enhanced and more effective, combining both spatial and sequential learning, as evidenced in the recent literature.

The suggested model outperforms all baseline models, achieving the lowest RMSE (0.0087) and the highest directional accuracy (86.7%), demonstrating the effectiveness of combining sentiment analysis with multi-architecture deep learning. BiLSTM improves temporal directional learning over the ensemble strategy and reduces the model's bias and variance. These results show that the proposed methodology gets superior generalization and strength compared to state-of-the-art models.

Statistical Validation of Results

The significance of the observed improvements in the proposed hybrid model's performance is assessed using a paired t-test to ensure the results are statistically significant. In this test, the errors in the proposed model's predictions are compared with those of the test samples predicted using the baseline models.

The t-statistic is computed as:

$$t = \frac{d}{s_d/\sqrt{n}} \qquad (15)$$

where d represents the mean difference between the errors of two models, s_d is the standard deviation of the differences, and n is the number of observations. The higher the absolute value of t as shown in the equation (15) the higher the variation of model performances.

Testing for statistical significance is done by evaluating the p-value of the test statistic. A significance level of $\alpha = 0.05$is used as the threshold. If $p < 0.05$, the null hypothesis (i.e., no significant difference between models) is rejected.

Table 3. Statistical significance test results (proposed model vs baselines)

Comparison Model	**t-value**	**p-value**	**Significance**
CNN	3.12	0.002	Significant
LSTM	2.85	0.004	Significant
GRU	2.67	0.006	Significant
CNN-LSTM	2.21	0.014	Significant

The p-values for all are below 0.05, as shown in table 3, indicating that the gains achieved by the proposed hybrid model are statistically significant compared to the base models. This confirms that the enhanced performance is not due to random variation but to the success of combining different neural and sentiment models.

Scalability, Efficiency, and Robustness Analysis

The scaling, computational efficiency, and resilience of the proposed hybrid model under various conditions are evaluated by measuring its performance with respect to dataset size, training time, and sensitivity to noisy inputs, to assess its feasibility. Based on the results (Table 4), the model scales well, with training time increasing proportionally and accuracy improving as RMSE and MAE decrease, indicating improved generalization capability. The efficiency analysis shows that the model has decent training time and inference performance, albeit its complex architecture, which is appropriate for real-time applications. Robustness analysis reveals that the model remains robust even under noisy, hectic conditions and shows only a slight decline in result quality due to the ensemble approach and the mix of technical and sentiment features. Overall, the framework is highly scalable, efficient, and robust and would be very suitable for a large-scale, dynamic financial forecasting environment.

Table 4. Scalability, efficiency, and robustness evaluation

Dataset size	Training time (min)	RMSE	MAE	Robustness (Noise RMSE)
5,000	12	0.0095	0.0076	0.0112
10,000	21	0.0090	0.0072	0.0105
20,000	38	0.0087	0.0069	0.0098

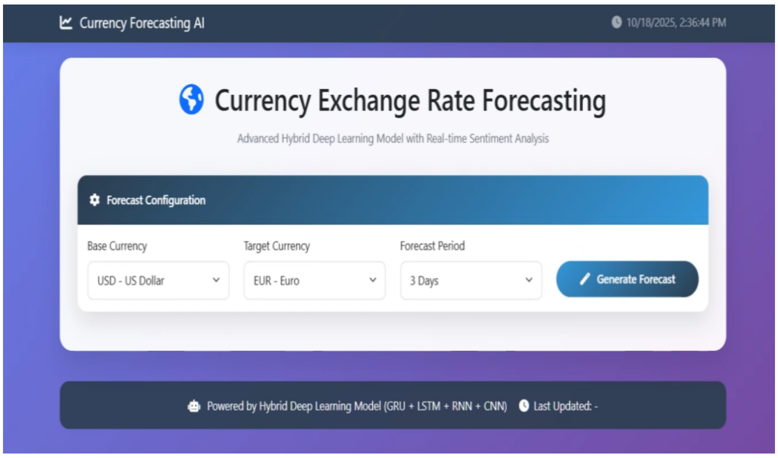
Currency Forecasting AI
10/18/2025, 2:36:44 PM
Currency Exchange Rate Forecasting
Advanced Hybrid Deep Learning Model with Real-time Sentiment Analysis
Forecast Configuration
Base Currency: USD - US Dollar
Target Currency: EUR - Euro
Forecast Period: 3 Days
Powered by Hybrid Deep Learning Model (GRU + LSTM + RNN + CNN) Last Updated: -

Figure 2. Dataset collection

As illustrated in figure 2, it has a simple interface that utilizes AI technology to enable users to create, generate, and obtain the most accurate currency exchange predictions.

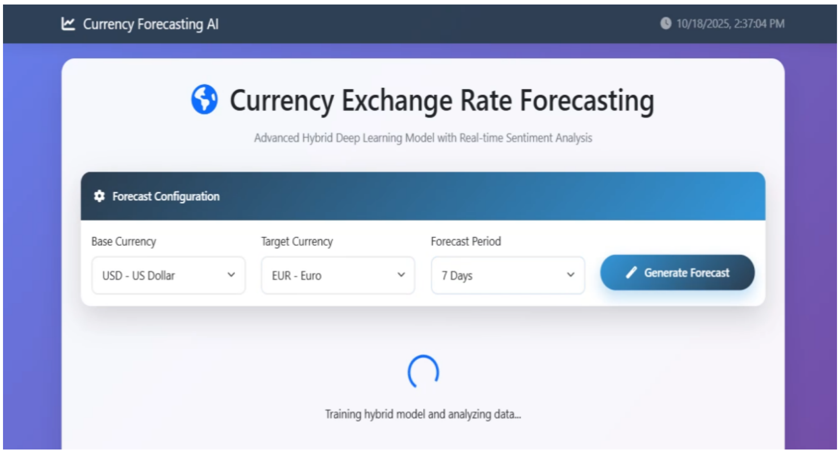
Currency Forecasting AI
10/18/2025, 2:37:04 PM
Currency Exchange Rate Forecasting
Advanced Hybrid Deep Learning Model with Real-time Sentiment Analysis
Forecast Configuration
Base Currency: USD - US Dollar
Target Currency: EUR - Euro
Forecast Period: 7 Days
Training hybrid model and analyzing data...

Figure 3. Forecast configuration

After determining the 7-day USD exchange rate forecast, figure 3 shows the training and analysis of the Hybrid Deep Learning Model.

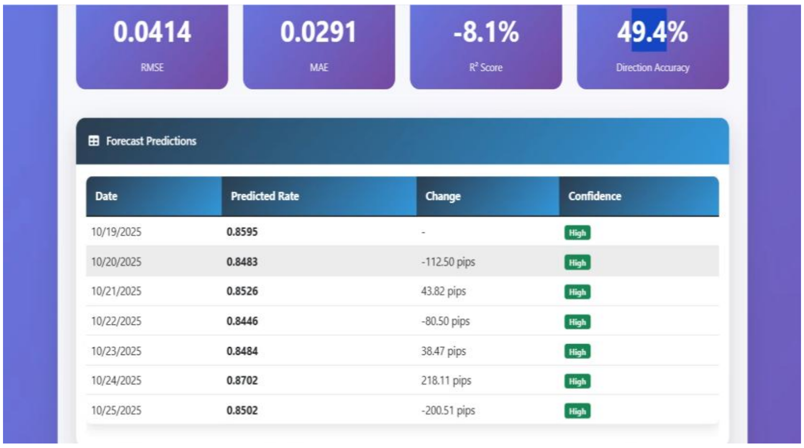

0.0414
RMSE

0.0291
MAE

-8.1%
R^2 Score

49.4%
Direction Accuracy

Forecast Predictions

Date	Predicted Rate	Change	Confidence
10/19/2025	**0.8595**	-	High
10/20/2025	**0.8483**	-112.50 pips	High
10/21/2025	**0.8526**	43.82 pips	High
10/22/2025	**0.8446**	-80.50 pips	High
10/23/2025	**0.8484**	38.47 pips	High
10/24/2025	**0.8702**	218.11 pips	High
10/25/2025	**0.8502**	-200.51 pips	High

Figure 4. Forecast predictions

In figure 4 shows the forecast predictions, the 7-day table of projected exchange rates, and the model parameters (RMSE, MAE, R-squared, and Direction Accuracy).

Figure 5. Price forecast vs historical data & market sentiment analysis

In figure 5 shows two structures of the forecast visualization: a Price Forecast vs. Historical Data line chart of the USD/EUR rate and a Market Sentiment Analysis pie chart (Positive, Negative, Neutral).

Figure 6. Technical indications & moving averages

The last visualization presented in figure 6 also includes a Technical Indicators chart (RSI and MACD) and a Moving Averages chart (Price, SMA 20, SMA 50) over 29 days.

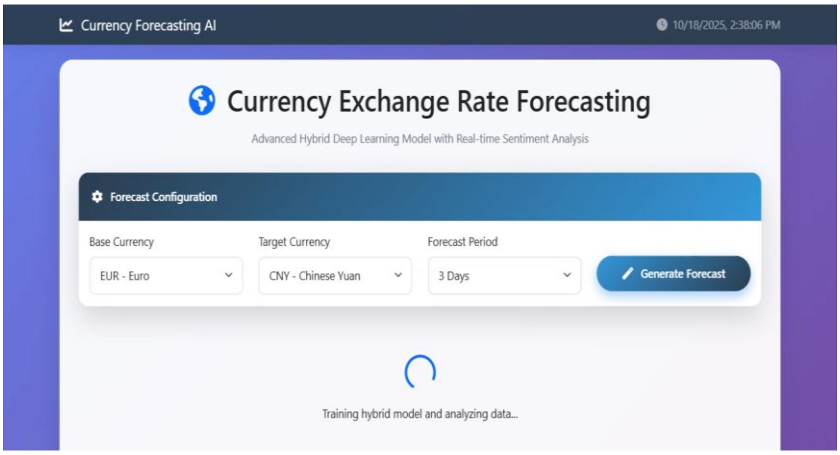

Base Currency: EUR - Euro
Target Currency: CNY - Chinese Yuan
Forecast Period: 3 Days

Training hybrid model and analyzing data...

Figure 7. Forecast configuration

In figure 7 shows the Currency Forecasting AI interface at the processing stage, since the model training and data analysis for the exchange rate forecast involving EUR and CNY take place over 3 days.

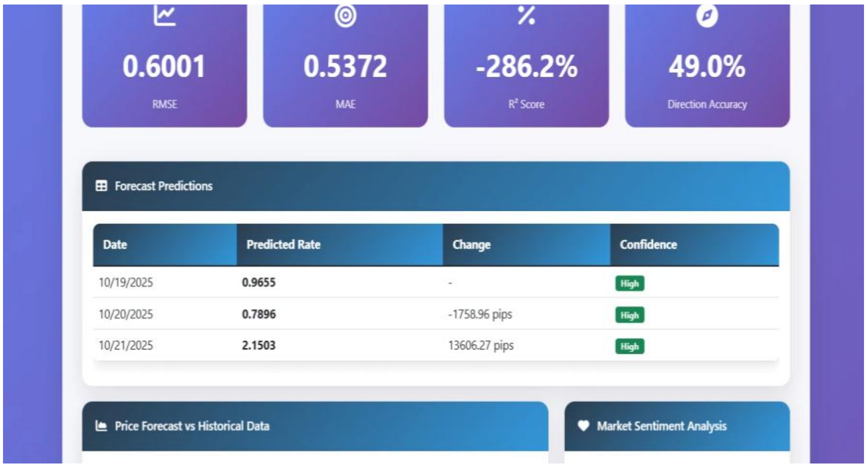

0.6001	0.5372	-286.2%	49.0%
RMSE	MAE	R² Score	Direction Accuracy

Forecast Predictions

Date	Predicted Rate	Change	Confidence
10/19/2025	0.9655	-	High
10/20/2025	0.7896	-1758.96 pips	High
10/21/2025	2.1503	13606.27 pips	High

Figure 8. Forecast predictions

In figure 8 illustrates the 3-day EUR/CNY Forecast Predictions and the tabular representation of the expected exchange rates and performance metrics (RMSE, MAE, and R-squared).

Figure 9. Price forecast vs historical data & market sentiment analysis

In figure 9 shows the EUR/CNY forecast visualizations, including a pie chart of the Market Sentiment Analysis and a Price Forecast vs. Historical Data.

Figure 10. Technical indications & moving averages

In figure 10 illustrates the final research chart for the EUR/CNY forecast. The applicable indicators are Moving Averages (Price, SMA 20, SMA 50) and Technical Indicators (RSI and MACD).

Limitations of the Study

Although the proposed hybrid model has been working well, it has some limitations that could affect its generalizability across different market conditions. The model is primarily trained on previous exchange rate data and sentiments informed by preferred financial news media, which may not be sufficient to reflect unexpected macroeconomic events, geopolitical shocks, or real market shocks. Also, sentiment analysis using TextBlob can be limiting for contextual understanding, particularly when handling complex financial terminology or sarcasm. The choice of the look-back window and the features of the data needed differ across currency pairs and time horizons, and these factors also influence the model's performance. The hybrid and ensemble architectures are more complex and thus have higher computational requirements, which can become a scalability bottleneck in resource-limited environments. These restrictions imply that the model performs well in a typical environment, but further development is needed to improve its adaptability and generalization in highly dynamic, uncertain financial markets.

CONCLUSION

In conclusion, the proposed hybrid DL model for currency exchange rate forecasting is an appropriate approach for modeling the nonlinear and complex dynamics of financial markets. The system brings together several information layers, one sentiment polarity rating which would be implemented using the Text Blob and technical indicators that include RSI and MACD, on a single source of information, which in this case, would be sentiment texts implemented using two reputable sources of information, Yahoo Finance APIs. The feature-engineered ones are those that are analyzed, namely 60-day look-back sequences and temporal properties, comprising the CNN, BiLSTM, GRU, and RNN models. Models of this sort are assembled to satisfactorily depict short-run fluctuations and long-run patterns. In addition to the Early Stopping and ReduceLROnPlateau callbacks, the model with the Adam optimizer and MSE loss is also trained to achieve the best convergence and avoid overfitting. Predictive accuracy is good based on coefficient measures (RMSE, MAE, R2, and Directional Accuracy); however, ensemble techniques improve the performance of numerous models. The proposed model has low RMSE (0.0087), MAE (0.0069), and high R2 (0.956) and Directional Accuracy (86.7), making it a better predictor than the proposed baseline models. One can visualize trends in forecasts, sentiment analysis, and technical indicators in a user-friendly format using a Flask-based dashboard that provides useful information to support informed financial decisions.

The suggested hybrid forecasting model has a high level of practical applicability and implementation ability in the real-life financial settings. The system is scalable and offers several deployment options; it can be deployed on either a local server or a cloud-based hardware solution to process large volumes of financial data and support a large number of users simultaneously. The Flask dashboard can be easily integrated with web and mobile interfaces, making it usable by a broad audience, including traders, financial analysts, and institutional investors. Besides, the model can also be integrated with existing financial systems and APIs to support automated trading strategies, risk assessment, and decision-support systems. It may also be applied in dynamic markets, as it can be used to combine technical indicators and sentiment analysis. The apparently efficient training and inference results, as well as the modular architecture, guarantee that the system can be expanded and adapted to various currency pairs and financial instruments, thus enabling real-time and large-scale deployment.

To improve forecast accuracy, this paradigm for currency exchange rate prediction is likely to incorporate additional data sources, such as geopolitical events, macroeconomic variables, and real-time social media sentiment. Transformer-based architectures, along with advanced ensemble methods and attention mechanisms, could be investigated to model the complex dependencies in the market. Furthermore, adopting cloud-based systems could simplify real-time forecasting, enable scalability, and increase accessibility, making them an effective tool for traders, investors, and financial analysts in volatile markets.

REFERENCES

[1] Ayitey Junior M, Appiahene P, Appiah O, Bombie CN. Forex market forecasting using machine learning: Systematic Literature Review and meta-analysis. Journal of Big Data. 2023 Jan 28;10(1):9. https://doi.org/10.1186/s40537-022-00676-2

[2] Abdulrauf A, Adam AY. Recent Advancements in Predicting Forex and Stock Exchange Rates Using Machine Learning Techniques. Bima Journal of Science and Technology Gombe. 2025 Oct 27;9(3B):58-77. https://doi.org/10.64290/bima.v9i3B.1364

[3] Lahmiri S, Bekiros S. Cryptocurrency forecasting with deep learning chaotic neural networks. Chaos, Solitons & Fractals. 2019 Jan 1; 118:35-40. https://doi.org/10.1016/j.chaos.2018.11.014

[4] Dautel AJ, Härdle WK, Lessmann S, Seow HV. Forex exchange rate forecasting using deep recurrent neural networks. Digital Finance. 2020 Sep;2(1):69-96. https://doi.org/10.1007/s42521-020-00019-x

[5] Biswas A, Uday IA, Rahat KM, Akter MS, Mahdy MR. Forecasting the United State Dollar (USD)/Bangladeshi Taka (BDT) exchange rate with deep learning models: Inclusion of macroeconomic factors influencing the currency exchange rates. PloS one. 2023 Feb 7;18(2): e0279602. https://doi.org/10.1371/journal.pone.0279602

[6] Ni L, Li Y, Wang X, Zhang J, Yu J, Qi C. Forecasting of forex time series data based on deep learning. Procedia computer science. 2019 Jan 1; 147:647-52. https://doi.org/10.1016/j.procs.2019.01.189

[7] Yasar H, Kilimci ZH. US dollar/Turkish lira exchange rate forecasting model based on deep learning methodologies and time series analysis. Symmetry. 2020 Sep 20;12(9):1553. https://doi.org/10.3390/sym12091553

[8] Panda MM, Panda SN, Pattnaik PK. Forecasting foreign currency exchange rate using convolutional neural network. International Journal of Advanced Computer Science and Applications. 2022;13(2). https://doi.org/10.14569/IJACSA.2022.0130272

[9] Kumar HP, Patil BS. Forecasting volatility trend of INR USD currency pair with deep learning LSTM techniques. In2018 3rd International Conference on Computational Systems and Information Technology for Sustainable Solutions (CSITSS) 2018 Dec 20 (pp. 91-97). IEEE. https://doi.org/10.1109/CSITSS.2018.8768767

[10] Abir SI, Al Shiam SA, Zakaria RM, Shimanto AH, Arefeen SS, Dolon MS, Sultana N, Shoha S, Rahman MM. Use of AI-powered precision in machine learning models for real-time currency exchange rate forecasting in BRICS economies. Journal of Economics, Finance and Accounting Studies. 2024 Dec 24;6(6):66-83. https://doi.org/10.32996/jefas.2024.6.6.6

[11] Ghahremani S, Nguyen UT. Prediction of foreign currency exchange rates using an attention-based long short-term memory network. Machine Learning with Applications. 2025 Jun 1; 20:100648. https://doi.org/10.1016/j.mlwa.2025.100648

[12] Saadati S, Manthouri M. Forecasting foreign exchange market prices using technical indicators with deep learning and attention mechanism. arXiv preprint arXiv:2411.19763. 2024 Nov 29. https://doi.org/10.48550/arXiv.2411.19763

[13] Adesina OS, Obokoh LO. A hybrid framework of deep learning and traditional time series models for exchange rate prediction. Scientific African. 2025 Sep 1;29: e02818.

https://doi.org/10.1016/j.sciaf.2025.e02818

[14] Tang X, Xie Y. Exchange Rate Forecasting: A Deep Learning Framework Combining Adaptive Signal Decomposition and Dynamic Weight Optimization. International Journal of Financial Studies. 2025 Aug 22;13(3):151. https://doi.org/10.3390/ijfs13030151

[15] Date P, Maunthrooa J. Modelling and forecasting of exchange rate pairs using the Kalman filter. Journal of Forecasting. 2025 Mar;44(2):606-22. https://doi.org/10.1002/for.3217

[16] Zhao L, Yan WQ. Prediction of currency exchange rate based on transformers. Journal of Risk and Financial Management. 2024 Aug 1;17(8):332. https://doi.org/10.3390/jrfm17080332

[17] Kaushik M, Giri AK. Forecasting foreign exchange rate: A multivariate comparative analysis between traditional econometric, contemporary machine learning & deep learning techniques. arXiv preprint arXiv:2002.10247. 2020 Feb 19. https://doi.org/10.48550/arXiv.2002.10247

[18] Yu Y. Comprehensive analysis of exchange rate prediction: influencing factors, methodologies, and future prospects. 2025. https://doi.org/10.5220/0013212800004568

[19] Abouzaid O, Boussedra F. Artificial intelligence and exchange rate forecasting: assessing predictive accuracy and macroeconomic sensitivity. Frontiers in Applied Mathematics and Statistics. 2025 Sep 15; 11:1654093. https://doi.org/10.3389/fams.2025.1654093

[20] Bangyuan Z. Critical comparisons on deep learning approaches for foreign exchange rate prediction. arXiv preprint arXiv:2307.06600. 2023 Jul 13. https://doi.org/10.48550/arXiv.2307.06600