“Decoding currency dynamics: A multiscale machine learning approach integrating economic indicators, ESG, and investor sentiment”

AUTHORS	Sougata Banerjee

ARTICLE INFO	Sougata Banerjee (2025). Decoding currency dynamics: A multiscale machine learning approach integrating economic indicators, ESG, and investor sentiment. *Investment Management and Financial Innovations*, *22*(3), 27-48. doi:10.21511/imfi.22(3).2025.03
DOI	http://dx.doi.org/10.21511/imfi.22(3).2025.03
RELEASED ON	Friday, 04 July 2025
RECEIVED ON	Sunday, 01 December 2024
ACCEPTED ON	Monday, 23 June 2025
LICENSE	This work is licensed under a Creative Commons Attribution 4.0 International License
JOURNAL	"Investment Management and Financial Innovations"
ISSN PRINT	1810-4967
ISSN ONLINE	1812-9358
PUBLISHER	LLC “Consulting Publishing Company “Business Perspectives”
FOUNDER	LLC “Consulting Publishing Company “Business Perspectives”

NUMBER OF REFERENCES	NUMBER OF FIGURES	NUMBER OF TABLES
44	17	5

BUSINESS PERSPECTIVES

LLC "CPC "Business Perspectives"
Hryhorii Skovoroda lane, 10,
Sumy, 40022, Ukraine
www.businessperspectives.org

Sougata Banerjee (India)

DECODING CURRENCY DYNAMICS: A MULTISCALE MACHINE LEARNING APPROACH INTEGRATING ECONOMIC INDICATORS, ESG, AND INVESTOR SENTIMENT

Abstract

The foreign exchange market, characterized by high volatility and economic significance, requires accurate predictive models. This study investigates the application of the Temporal Fusion Transformer (TFT), enhanced with Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN), for forecasting major foreign exchange (forex) currency pairs: USD/EUR, USD/JPY, USD/CNY, USD/AUD, and USD/INR. The proposed framework integrates a wide range of economic indicators, which include interest rate differentials, GDP growth, and trade balances, alongside investor sentiment derived from Twitter and ESG-related news sentiment. By addressing the non-linear, multiscale nature of forex time series, the CEEMDAN-TFT model facilitates improved signal decomposition and interpretability. Empirical results indicate that the model demonstrates competitive performance across all five currency pairs, with the USD/EUR pair exhibiting the highest predictive accuracy. Other pairs, exhibiting good predictive accuracy, include USD/JPY and USD/CNY, underscoring the model's adaptability to varying economic contexts. Performance is assessed using multiple error metrics, and the model is benchmarked against standard neural network approaches (MLP, RNN, LSTM, CNN). Variable importance analysis highlights the dynamic influence of interest rates, investor sentiment, and ESG factors across different market regimes. This study provides empirical evidence that including ESG and investor sentiment can improve the accuracy of currency forecasting models. This study provides guidance and a framework for informed decision-making for traders, analysts, and policymakers.

Received on: 1st of December, 2024
Accepted on: 23rd of June, 2025
Published on: 4th of July, 2025

Sougata Banerjee, Ph.D., Indian Institute of Management Ranchi, Finance Department, Indian Institute of Management, India.

Keywords forex forecasting, Temporal Fusion Transformer, machine learning, investor sentiment, ESG factor; CEEMDAN decomposition, multiscale decomposition, sentiment analysis

JEL Classification C53, F31, G17, Q56

INTRODUCTION

The foreign exchange (forex) market, distinguished by its immense trading volume and inherent volatility, occupies a pivotal position within the global financial system. The foreign exchange (forex) market is characterized by high volatility and complexity, influenced by factors such as economic fundamentals, market sentiment, geopolitical developments, and speculative trading. These dynamics pose substantial forecasting challenges (Menkhoff & Taylor, 2007). As a result, developing accurate exchange rate prediction models remains a key concern for investors, policymakers, and financial institutions (Embrechts, 1995). Despite their utility, traditional econometric models such as ARIMA and GARCH, while effective for specific conditions, often fail to capture the non-linear dynamics and real-time fluctuations characteristic inherent in financial time series (Engle, 2001;

Conflict of interest statement:
Author(s) reported no conflict of interest

Panda et al., 2022). These limitations are exacerbated by the shift to floating exchange rate regimes following the collapse of the Bretton Woods system (MacDonald, 2007), aligning with the Efficient Market Hypothesis (EMH), which suggests that exchange rates follow a random walk and are inherently unpredictable (Fama, 1970). Despite this, inefficiencies and patterns in forex markets indicate potential for improved modeling techniques that can leverage these complexities.

Recent advancements in machine learning (ML) provide powerful tools to address these challenges. ML models such as RNNs and LSTMs have demonstrated their ability to process high-dimensional and sequential data, offering a pathway to capture the intricate dependencies in forex time series (Hornik,1991; Weigend et al., 1992). However, their limitations in interpretability and long-term dependency modeling remain significant obstacles (Huy et al., 2022; Zhang et al., 2022). Complementary approaches like the Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) offer a promising solution by breaking down time series data into meaningful components, isolating noise, and revealing trends across different scales (Torres et al., 2011; Abraham, 1982). Incorporating investor sentiment from platforms like Twitter and ESG news sentiment adds a critical qualitative dimension to forex forecasting, addressing the behavioral and sustainability factors that influence market movements (Mehta et al., 2021; Bose et al., 2021). These perspectives reflect an evolving interest in incorporating psychological and macroeconomic influences into financial modeling, extending beyond traditional indicator-based approaches. In this context, there remains a need for models that can flexibly accommodate the multiscale and dynamic nature of forex markets while integrating a combination of economic fundamentals, investor behavior, and sustainability-related signals. This study aims to develop a comprehensive and interpretable model that may contribute to the understanding of currency market behavior and support ongoing efforts to improve forecasting accuracy.

1. LITERATURE REVIEW

Time series forecasting in financial markets, especially forex markets, has long been a critical area of research due to the volatile and non-linear nature of exchange rate movements (Sharma et al., 2021). Traditional econometric models, including ARIMA and GARCH, have been extensively applied to model forex rates (Engle, 2001; Sharma et al., 2021); however, these models face significant limitations when dealing with non-linear dependencies and the stochastic nature of forex markets. Empirical studies have shown that such models often perform no better than the random walk model for out-of-sample forecasts, highlighting their inability to effectively model exchange rate dynamics (Meese & Rogoff, 1983; Alexander & Thomas, 1987). The adoption of machine learning (ML) techniques has provided an alternative approach to capturing complex patterns in financial data. Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and Convolutional Neural Networks (CNNs) have demonstrated their efficacy in modeling sequential data, including forex rates (Hornik, 1991; Weigend et al., 1992). Despite their success, these models often lack interpretability and struggle with long-term dependencies (Huy et al., 2022; Zhang et al., 2022).

The Temporal Fusion Transformer (TFT) addresses these limitations by combining attention mechanisms and interpretable features, making it particularly suitable for multi-horizon forecasting tasks (Lim et al., 2021). Advanced decomposition techniques, such as the Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN), further enhance forecasting accuracy by isolating noise and breaking down time series data into meaningful components (Torres et al., 2011; Cao et al., 2019). CEEMDAN's ability to handle non-stationary and multi-scale time series has proven particularly valuable in financial applications, where trends, seasonality, and high-frequency fluctuations coexist (Abraham, 1982; Mohammadipour & Boylan, 2012).

Simultaneously, investor sentiment and environmental, social, and governance (ESG) factors are gaining recognition as significant drivers of forex market movements (Chen et al., 2024; Steuer & Troger, 2022; Abdel Magid et al., 2023). Social me-

http://dx.doi.org/10.21511/imfi.22(3).2025.03

dia platforms like Twitter have been used to extract sentiment indicators, which reflect investor mood and influence financial decision-making (Mehta et al., 2021; Safari Bideskan et al., 2022). Similarly, ESG factors reflect broader societal trends and have been shown to impact financial markets, making them crucial considerations for holistic predictive models (Bose et al., 2021; Huang et al., 2024). Integrating these qualitative dimensions with traditional economic indicators, such as interest rate differentials and trade balances, offers a more comprehensive framework for forecasting currency movements (Jan & Gopalaswamy, 2019; Bhattacharya, 2012). Several recent studies have established beyond doubt the importance of multiscale and dynamic modeling in capturing the nuances of forex markets (Hong & Zu, 2024; Bekiros & Marcellino, 2013). Similarly, multiple studies have demonstrated that multiscale approaches, such as combining CEEMDAN with ML models, improve the accuracy in predicting exchange rates by addressing limitations of single-scale methods (Kourentzes et al., 2014; Zhang et al., 2017). Additionally, the use of attention-based architectures, such as the TFT, enables dynamic weighting of influential variables, allowing models to adapt to changing market conditions and enhance interpretability (Lim et al., 2021; Simpson, 2002).

The extant literature has numerous studies focusing on the application of advanced deep learning models to exchange rate forecasting. For instance, Zhang et al. (2022) used transformer-based architectures to model multi-horizon financial time series, and Yilmaz and Arabaci (2021) explored the application of deep learning models for predicting foreign exchange returns by including real economic indicators as inputs. However, the literature lacks a comprehensive framework for combining ESG, investor sentiment, and advanced machine learning to forecast exchange rates. This study aims to address this gap by integrating economic indicators, investor sentiment, and ESG factors into a unified forecasting framework to predict forex rates for major currency pairs, leveraging the CEEMDAN and TFT models. Therefore, the following hypotheses are proposed:

H1: *The CEEMDAN-TFT model will outperform traditional neural network models (MLP, RNN, LSTM, CNN) in predictive accuracy for major currency pairs.*

H2: *Investor sentiment, quantified through Twitter data, significantly influences forex rate predictions, particularly during periods of market volatility.*

H3: *ESG factors, measured through sentiment scores, play a critical role in forecasting forex rates by capturing broader economic and societal trends.*

H4: *The integration of multi-scale decomposition techniques (CEEMDAN) with dynamic attention-based models (TFT) enhances interpretability and improves predictive performance across varying market conditions.*

2. METHODOLOGY

This study compiled a dataset that includes multiple economic variables as suggested in the existing literature, and a novel attempt at incorporating both ESG factors and investor sentiment as predictor variables. The dataset spans a period from January 2019 to March 2024, with daily frequency, excluding weekends and holidays. All data except investor sentiment are sourced from Bloomberg. The first set of variables includes the forex rates for five major internationally traded currency pairs. These pairs were selected based on their prominence in global financial markets and coverage in prior research. The USD/EUR exchange rate is included, as highlighted by Ito and Sato (2008), particularly in the context of carry trades and interest rate differentials. The USD/CNY exchange rate, studied by Huang et al. (2024) and Guo and Chen (2023), is included for its relevance to China's exchange rate policy and broader implications for global economic stability. The USD/AUD pair, analyzed by Chen et al. (2010), is incorporated to capture the relationship between commodity prices and exchange rates, especially in commodity-exporting countries like Australia. Finally, the USD/INR exchange rate is included as a reflection of the impact of macroeconomic variables and capital flows, as studied by Bhattacharya (2012) and Suman and Raju (2024).

In addition to the forex rates, the study included several other variables as per the existing study by Plakandaras et al. (2015), such as major stock indices from key global markets, including the U.S., Europe, China, Australia, India, and Japan. These indices provide a broad view of global equity market dynamics, which is crucial for understanding their influence on forex rates. As Plakandaras et al. (2015) suggested, the dataset further incorporates spot prices for ten precious and non-precious metals, acknowledging the impact of commodity markets on currency movements. Additionally, the prices of 19 key commodities, including crude oil, are included to account for the influence of global commodity price fluctuations on exchange rates. To capture the temporal dynamics and trends in the forex market, the study calculates moving averages over 3, 5, 10, and 30 days for all included exchange rates. These moving averages serve as important technical indicators, smoothing out short-term volatility and highlighting longer-term trends. The dataset also includes four major macroeconomic variables: inflation rate differentials, interest rate differentials of various maturities, interest rate spreads between commercial paper, the federal funds rate, and the effective federal funds rate, GDP growth rate differentials, and trade balances between the countries involved in the currency pairs. These macroeconomic factors are essential for understanding the underlying economic conditions that drive forex rate movements.

Furthermore, this study includes ESG sentiment scores calculated from ESG-related news sourced from Bloomberg to quantify the market's response to ESG-related news and its influence on forex rates. This novel inclusion reflects the growing importance of sustainability factors in financial markets and their potential impact on currency values. Finally, investor sentiment is captured through scores derived from raw tweet text collected via the Twitter API following Safari Bideskan et al. (2022). These texts are analyzed using the VADER (Valence Aware Dictionary and sEntiment Reasoner) lexicon-based NLP algorithm to quantify sentiment in the market (Bose et al., 2021). This approach allows for real-time assessment of investor mood, which is increasingly recognized as a significant factor in financial market behavior (Mehta et al., 2021). These efforts resulted in a dataset comprising 76 variables as shown in Table 1 (adapted from Plakandaras et al., 2015), the variables encompass a range of economic and sentiment indicators compiled from Bloomberg, X (formerly Twitter) API, and informed by previous studies (Plakandaras et al., 2015; Bhattacharya, 2012).

Table 1. Variables employed in the forecasting model

Source: Adapted from Plakandaras et al. (2015).

Commodities
Crude oil
Cotton
Lumber
Cocoa
Coffee
Orange juice
Sugar
Corn
Wheat
Oats
Rough rice
Soybean meal
Soybean oil
Soybeans
Feeder cattle
Lean hogs
Live cattle
Pork bellies
Iron ore
Metals
Gold
Copper
Palladium
Platinum
Silver
Aluminum
Zinc
Nickel
Lead
Tin
Technical analysis variables
5-day index
Moving average 3-day
Moving average 5-day
Moving average 10-day
Moving average 30-day
Exchange rates
USD/EUR
USD/JPY
USD/CYN
USD/AUD
USD/INR
Interest rates
T-bill 6 months
T-bill 10 years
Spread MLP-EURIBOR 3M
Spread MLR-Eonia
Spread FF-CP
Spread FF-EFF

http://dx.doi.org/10.21511/imfi.22(3).2025.03

Table 1 (cont.). Variables employed in the forecasting model

Interest rates
EONIA
EURIBOR 1 week
ECB Interest rate
EURIBOR 1 month
FED rate
Stock indices
Dow Jones
Nasdaq 100
S&P 500
SSE
S&P/ASX 200
NIFTY 50
BSE Sensex
Euro Stoxx 50
Macroeconomic Variables for all countries
Consumer price index
Productivity index
Gross domestic product
Trade balance
Unemployment rate
Central bank discount rate
Long-term interest rate
Short-term interest rate
Aggregate money M3
Public debt
Deficit/surplus of the government budget
Sentiment
Positive IS score
Negative IS score
Composite IS score
E – score
S – score
G – score
ESG – composite score

2.1. Decomposition method

Forex rates are inherently volatile, as observed by MacDonald (2007). Kourentzes et al. (2014) suggest that instead of eliminating the volatile components of the time series, forecasting at different sampling frequencies (e.g., daily, monthly, yearly) and combining these forecasts might be more effective. High-frequency data captures short-run dynamics, while lower-frequency data reveals long-term characteristics like trends and seasonality. This multiscale approach helps better identify different aspects of the time series, improving forecasting accuracy and reducing bias (Abraham, 1982; Mohammadipour & Boylan, 2012). This study utilized the CEEMDAN technique to analyze the multiscale nature and nonlinearity of the forex rate series. CEEMDAN was chosen for its ability to overcome the mode-mixing limitations of earlier decomposition methods like EMD and EEMD, ensuring a more accurate and interpretable decomposition of time series components (Torres et al., 2011). Its ability to decompose non-stationary and nonlinear time series into intrinsic mode functions (IMFs) allows for extracting meaningful temporal patterns across different scales. This is particularly relevant in financial applications like forex forecasting, where non-stationarity and volatility are prevalent (Zhang et al., 2017). This study adopted a similar methodology for applying CEEMDAN as the existing study by Cao et al. (2019), and the decomposition was conducted using the PyEMD library in Python.

2.2. Forecasting model: Temporal Fusion Transformer (TFT)

The Temporal Fusion Transformer (TFT), introduced by Lim et al. (2021), is particularly suited for forex forecasting due to its ability to integrate diverse data sources, such as economic indicators, market sentiment, and external factors. TFT offers interpretability, enabling insights into the drivers of forex movements, and is robust to missing data, making it reliable for real-world applications. Its dynamic processing of covariates allows it to adjust variable importance as market conditions change, capturing both short- and long-term dependencies through its combination of recurrent networks, attention mechanisms, and interpretable modeling (Huy et al., 2022; Zhang et al., 2022). Below is a detailed explanation of the TFT architecture.

TFT takes as input a multivariate time series:

$$X = \{x_t\}_{t=1}^{T}, \tag{1}$$

where X_t Î R^d represents the input features at time t, the inputs can be categorized into

(i) Static covariates (S Î R^k), which are features that do not change over time, such as the identity of the currency pair.

(ii) Known inputs $\left(X^{known} = \{x_t^{known}\}_{t=1}^{T}\right)$, features known ahead of time, like macroeconomic indicators or scheduled events.

(iii) Observed inputs $\left(X^{obs} = \left\{ x_t^{obs} \right\}_{t=1}^{T} \right)$, features observed in real-time, such as historical forex rates or sentiment scores.

A key feature of TFT is the variable selection network, which dynamically selects relevant input features at each time step. The selection is achieved through a gating mechanism that weighs the importance of each feature. For a given input x_t, the variable selection can be expressed as:

$$V_t = \sigma\left(W_v x_t + b_v \right), \quad (2)$$

where W_v and b_v are trainable weights and biases, and σ(.) denotes the sigmoid activation function, which ensures the gating values are between 0 and 1. V_t Î R^d represents the selection vector that scales the input features.

Static covariates S are processed separately through a static covariate encoder, which outputs a static context vector C_s. This context vector influences the entire sequence by modulating other parts of the network. The static context vector is given by:

$$C_s = ReLU\left(W_s S + b_s \right), \quad (3)$$

where W_s and B_s are trainable weights and biases, respectively. *ReLU*(.) denotes the Rectified Linear Unit activation function, which introduces non-linearity. The selected variables V_t are passed through a Long Short-Term Memory (LSTM) network, which is well-suited for capturing temporal dependencies in the data. The LSTM update equations are as follows:

$$i_t = \sigma\left(W_i v_t + U_i h_{t-1} + b_i \right), \quad (4)$$

$$f_t = \sigma\left(W_f v_t + U_f h_{t-1} + b_f \right), \quad (5)$$

$$o_t = \sigma\left(W_o v_t + U_o h_{t-1} + b_o \right), \quad (6)$$

$$g_t = \tanh\left(W_g v_t + U_g h_{t-1} + b_g \right), \quad (7)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot g_t, \quad (8)$$

$$h_t = o_t \odot \tanh\left(c_t \right), \quad (9)$$

where i_t, f_t, o_t are the input, forget, and output gates, respectively. C_t and h_t represent the cell state and hidden state at time *t*. Whereas ⊙ denotes element-wise multiplication, W_i, W_f, W_o, W_g, U_i, U_f, U_o, and U_g are trainable weight matrices, and b_i, b_f, b_o, and b_g are biases.

TFT employs a multi-head attention mechanism to allow the model to focus on different parts of the input sequence at different time steps. This attention mechanism enhances the model's ability to capture long-range dependencies and complex interactions between different time steps. The scaled dot-product attention for each head is given by:

$$Attention\left(Q, K, V \right) = soft\max\left(\frac{QK^T}{\sqrt{d_k}} \right) V, \quad (10)$$

where Q, K, V are the query, key, and value matrices, respectively, derived from the input. D_k is the dimensionality of the key vectors. Multi-head attention is an extension where the attention mechanism is applied in parallel across multiple heads:

$$MultiHead\left(Q,\ K,\ V \right) = Concat\left(head_1,\ head_2,\ \ldots,\ head_h \right) W_o, \quad (11)$$

where $head_i = Attention\left(QW_i^Q, KW_i^K, VW_i^V \right)$ and W_i^Q, W_i^K, W_i^V are the projection matrices for each head. While *h* is the number of heads, and W_o is the output weight matrix. The Gated Residual Network (GRN) is used within TFT to process static and dynamic covariates. The GRN helps the network to flexibly adjust the contribution of each input and enhances the model's stability. The GRN is defined as:

$$z = W_2 ReLU \cdot \left(W_1 x + b_1 \right) + b_2, \quad (12)$$

$$g = \sigma\left(W_g x + b_g \right), \quad (13)$$

$$y = g \odot z + \left(1 - g \right) \odot x, \quad (14)$$

where *x* is the input vector. W_1, W_2, and W_g are weight matrices, as b_1, b_2, and b_g are biases, and *g* is the gating vector controlling the flow of information. Finally, the outputs from the LSTM and attention mechanisms are combined and passed through a fully connected layer to generate the final prediction. The output prediction $\widehat{y_t}$ at time *t* is given by:

$$\widehat{y_t} = FC\left(h_t \oplus c_s \right), \quad (15)$$

http://dx.doi.org/10.21511/imfi.22(3).2025.03

where $\oplus$ denotes concatenation, and the $FC(.)$ is a fully connected layer that maps the concatenated vector to the output space. The model is typically trained using a mean squared error (MSE) loss function, defined as:

$$L = \frac{1}{T}\sum_{t=1}^{T}\left(\widehat{y_t} - y_t\right)^2, \tag{16}$$

where $\widehat{y_t}$ is the predicted value at time t, and y_t is the actual value at time t. The TFT architecture leverages LSTM networks for temporal processing, multi-head attention for capturing complex dependencies, and gating mechanisms for dynamic feature selection. This combination allows TFT to effectively model multivariate time series data, making it particularly well-suited for forecasting tasks in financial markets, such as predicting forex rates. The architecture's ability to integrate both static and dynamic covariates, along with its interpretability features, makes it a powerful tool for generating accurate and actionable forecasts in complex domains.

2.3. Model evaluation approach

To assess the predictive performance of the CEEMDAN-TFT model and compare it with traditional neural network models (MLP, RNN, LSTM, CNN), the following set of error metrics was employed:

Mean Squared Error (MSE): Measures the average of the squared differences between predicted and actual values. Lower values indicate better performance. In highly volatile forex markets, controlling large deviations is critical for risk-sensitive applications like hedging and reserve management.

$$MSE = \frac{1}{n}\sum_{t=1}^{n}(\widehat{y_t} - y_t)^2. \tag{17}$$

Mean Absolute Error (MAE): MAE is a scale-dependent metric that gives a direct interpretation of forecast accuracy in the same units as the dependent variable. It is less sensitive to outliers compared to MSE, providing a robust measure of average prediction error for currency rate changes.

$$MAE = \frac{1}{n}\sum_{t=1}^{n}\left|\widehat{y_t} - y_t\right|. \tag{18}$$

Root Mean Squared Error (RMSE): RMSE is useful in financial contexts because it expresses errors in the same units as forex rates and disproportionately penalizes large errors. This makes it highly relevant for capturing sudden market fluctuations and currency shocks.

$$RMSE = \sqrt{MSE}. \tag{19}$$

R-squared (R^2): Indicates the percentage of variance in the actual forex rate explained by the model. High R^2 values indicate a model's strong explanatory power.

$$R^2 = 1 - \frac{\sum_{t=1}^{n}\left(\widehat{Y_t} - y_t\right)^2}{\sum_{t=1}^{n}\left(y_t - \overline{y_t}\right)^2}. \tag{20}$$

Mean Absolute Percentage Error (MAPE): Expresses accuracy as a percentage, making it easy to compare across currency pairs with different value scales. It is widely used in economic and financial forecasting literature (De Myttenaere et al., 2016).

$$MAPE = \frac{1}{n}\sum_{t=1}^{n}\left|\frac{\widehat{y_t} - y_t}{y_t}\right| \cdot 100. \tag{21}$$

Symmetric Mean Absolute Percentage Error (SMAPE): Normalizes MAPE to avoid inflated values when actual values are small. In forex contexts where daily returns can be near-zero, SMAPE offers a fairer error assessment.

$$SMAPE = \frac{100\%}{n}\sum_{t=1}^{n}\frac{\left|\widehat{y_t} - y_t\right|}{\left(\left|y_t\right| + \left|\widehat{y_t}\right|\right)/2}. \tag{22}$$

Mean Squared Logarithmic Error (MSLE): Penalizes under-predictions more than over-predictions. MSLE is particularly useful when currency returns or rates span different orders of magnitude, or when minimizing relative errors in log-transformed rates is prioritized. It dampens the effect of large outliers and emphasizes under-prediction.

$$MSLE = \frac{1}{n}\sum_{t=1}^{n}\left(\log\left(1+\widehat{y_t}\right) - \log\left(1+y_t\right)\right)^2. \tag{23}$$

These metrics were computed on both in-sample and out-of-sample forecasts to evaluate model fit and generalization capacity. All models were trained and validated on identical data splits. The

TFT model utilized early stopping based on validation loss, with a batch size of 64 and an initial learning rate of 0.001. Hyperparameters for all models were tuned using grid search to ensure fairness.

3. RESULT ANALYSIS AND DISCUSSION

Table 2 presents the summary statistics for the five selected currency exchange rates, their log-transformed values, and daily log returns. Variance analysis shows that raw exchange rate series have higher variance, especially in volatile pairs like USD/JPY (std. dev 15.8) and USD/INR (std. dev 4.63). Log transformations reduce variance, stabilizing the series, while return series (R) consistently exhibit very low, normalized variance. The skewness and kurtosis values are relatively low for USD/EUR and USD/AUD, suggesting that these series might exhibit linearity. In contrast, R(USD/JPY) and R(USD/CNY) show higher kurtosis values (6.98 and 4.36, respectively), indicating potential nonlinearity due to the presence of extreme values (fat tails). The relatively low skewness values across the log-transformed series suggest that these transformations help maintain linearity. The Jarque-Bera (JB) test results corroborate this finding. USD/EUR and log (USD/EUR) have p-values of 0.00 and 0.01, respectively, suggesting that these series are not normally distributed. High JB statistics and low p-values across the R(USD/JPY) and R(USD/INR) series also indicate significant deviations from normality, which is further supported by the extreme kurtosis values (6.98 for R(USD/JPY) and 3.69 for R(USD/INR)). The relatively lower skewness and kurtosis in the log-transformed series indicate some degree of normality improvement, but overall, most of the series reject the null hypothesis of normality. Figures 1 through 5 illustrate the daily rate movements for the selected currencies. To further analyze the stationarity of the series, unit root tests were conducted using the Augmented Dickey-Fuller (ADF), Phillips-Perron (PP), and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) tests. The results shown in Table 3 indicate R(USD/EUR), R(USD/JPY), R(USD/CNY), and R(USD/INR), the ADF and PP tests returned very low p-values (0.00), indicating rejection of the null hypothesis of a unit root; hence these series are stationary. The KPSS test generally supports the ADF and PP test results, except for the original and log-transformed series of USD/JPY and USD/INR, which show higher KPSS statistics with low p-values (indicating non-stationarity). The mixed results between the ADF, PP, and KPSS tests suggest that log transformations and differencing (returns) help achieve stationarity.

Table 2. Descriptive statistics

Variable Name	Mean	Std Dev	Min	Max	Skewness	Kurtosis	Jarque–Bera	JB p-value
USD/EUR	0.90	0.05	0.81	1.04	0.35	0.13	28.1	0.00
log (USD/EUR)	–0.05	0.02	–0.09	0.02	0.20	–0.07	9.3	0.01
R(USD/EUR)	0.00	0.00	–0.01	0.01	0.06	1.60	146.6	0.00
USD/JPY	120.9	15.86	102.4	151.7	0.65	–1.22	179.9	0.00
log (USD/JPY)	2.08	0.06	2.01	2.18	0.58	–1.32	175.8	0.00
R(USD/JPY)	0.00	0.00	–0.01	0.02	0.46	6.98	2821.5	0.00
USD/CNY	6.83	0.29	6.31	7.34	–0.21	–1.19	85.1	0.00
log (USD/CNY)	0.83	0.02	0.80	0.87	–0.25	–1.18	88.5	0.00
R(USD/CNY)	0.00	0.00	–0.01	0.01	0.10	4.36	1013.7	0.00
USD/AUD	1.44	0.08	1.26	1.74	0.10	–0.17	3.9	0.14
log (USD/AUD)	0.16	0.02	0.10	0.24	–0.05	–0.31	5.8	0.05
R(USD/AUD)	0.00	0.00	–0.02	0.01	–0.23	2.64	407.8	0.00
USD/INR	76.26	4.63	68.42	83.43	0.27	–1.29	104.3	0.00
log (USD/INR)	1.88	0.03	1.84	1.92	0.21	–1.27	95.9	0.00
R(USD/INR)	0.00	0.00	–0.01	0.01	–0.61	3.69	809.7	0.00

http://dx.doi.org/10.21511/imfi.22(3).2025.03

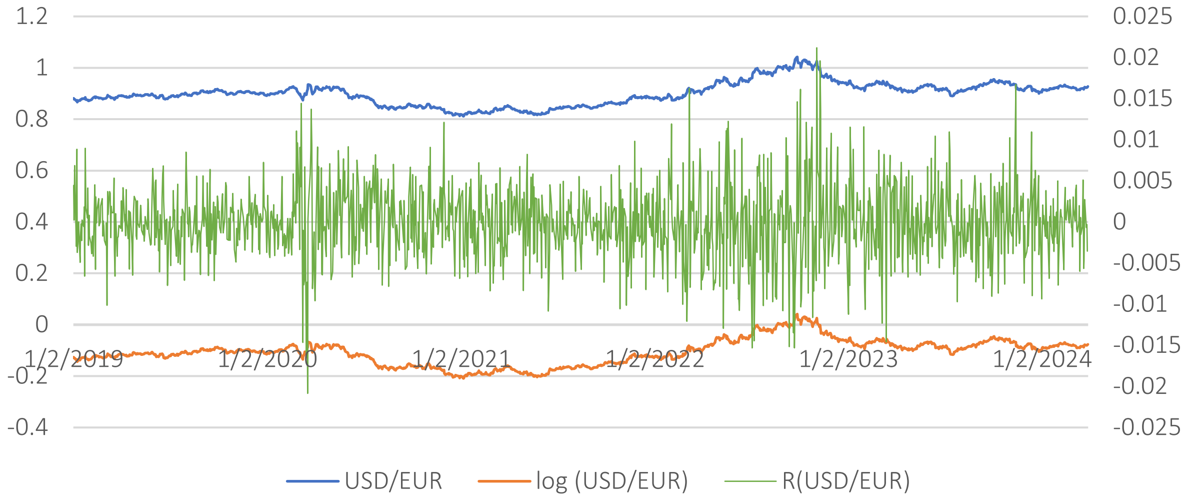

Figure 1. Daily rate movements of the USD and EUR pair

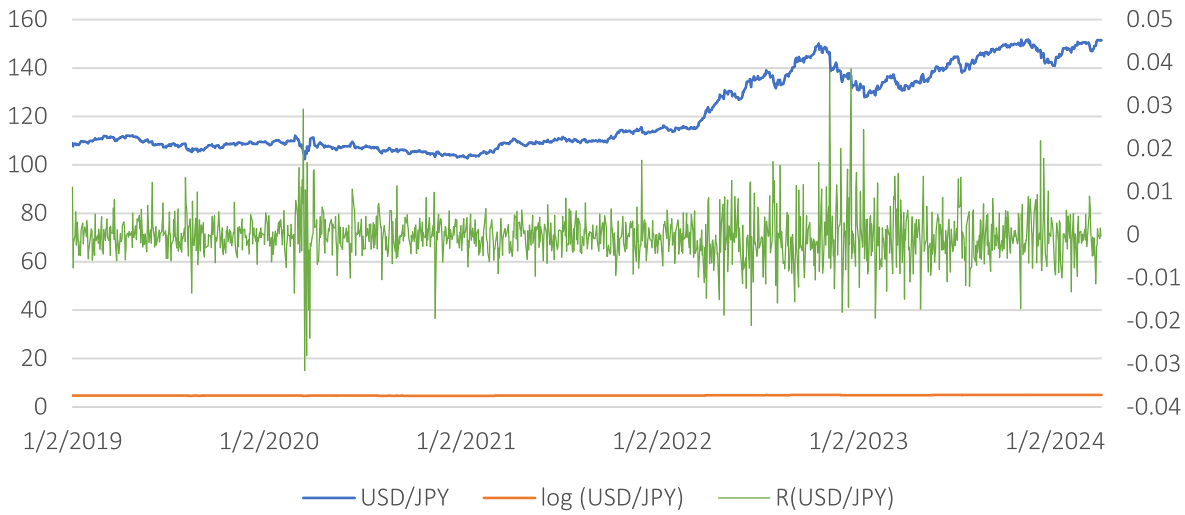

Figure 2. Daily rate movements of the USD and JPY pair

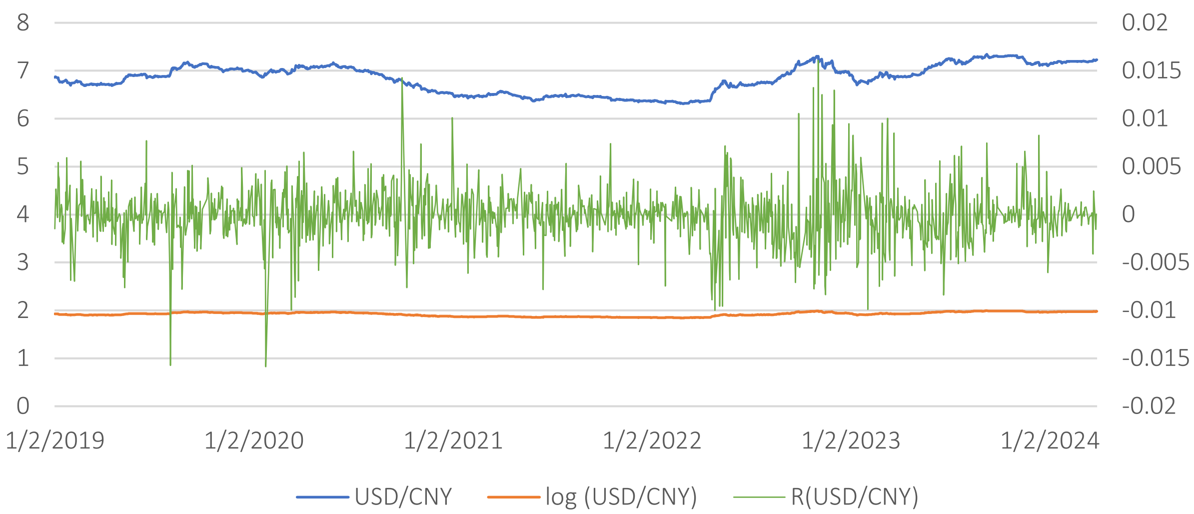

Figure 3. Daily rate movements of the USD and CNY pair

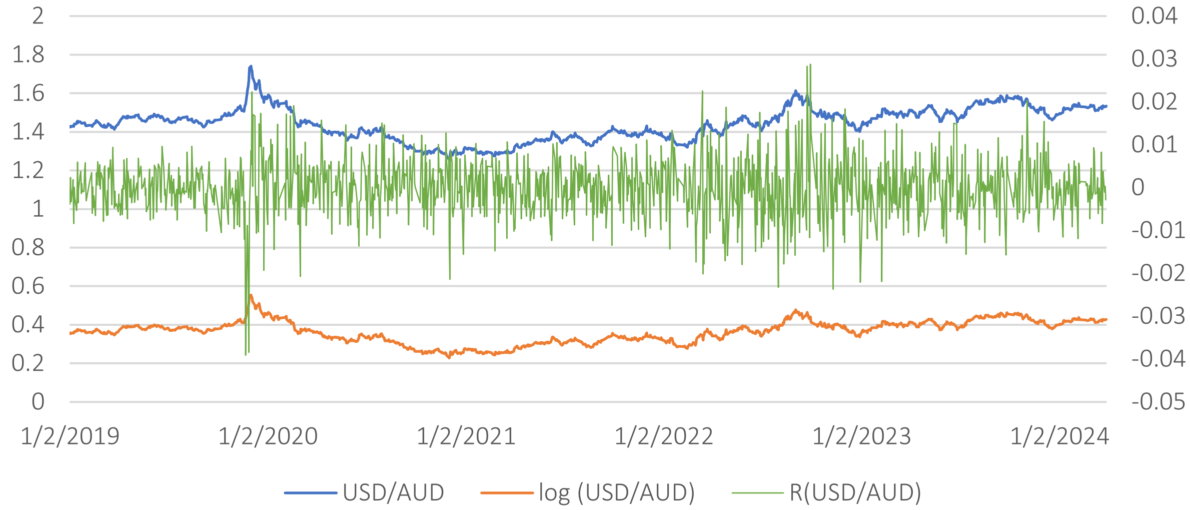

Figure 4. Daily rate movements of the USD and AUD pair

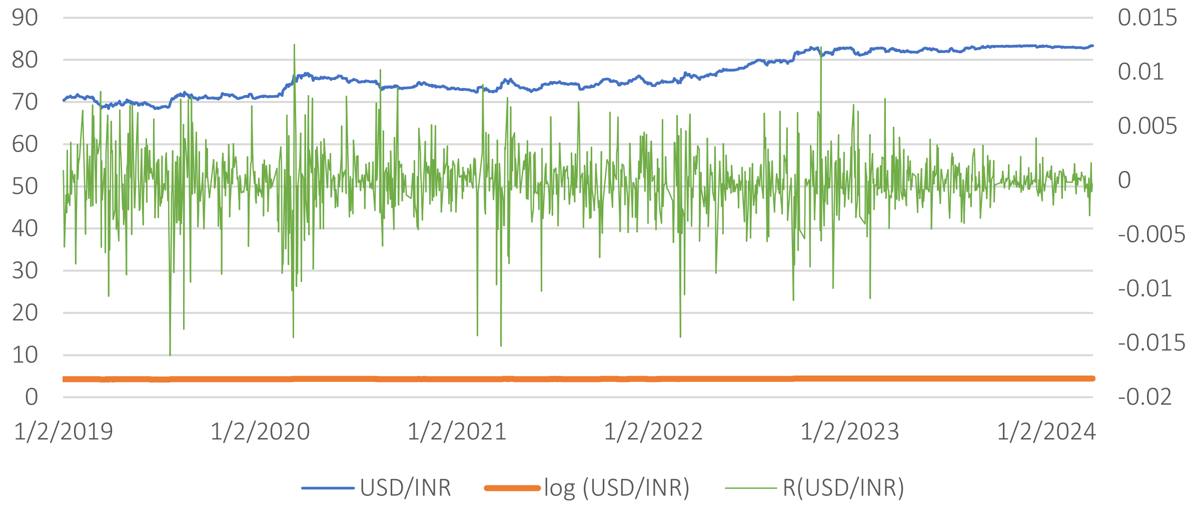

Figure 5. Daily rate movements of the USD and INR pair

Table 3. ADF, PP, and KPSS tests for a random walk

Variable	ADF Stat	ADF p-value	PP Stat	PP p-value	KPSS Stat	KPSS p-value
USD/EUR	–1.532	0.518	–1.632	0.466	1.987	0.010
log (USD/EUR)	–1.520	0.524	–1.611	0.477	1.983	0.010
R(USD/EUR)	–16.735	0.000	–36.218	0.000	0.074	0.100
USD/JPY	–1.776	0.392	–1.777	0.392	5.258	0.010
log (USD/JPY)	–1.677	0.443	–1.681	0.441	5.296	0.010
R(USD/JPY)	–37.205	0.000	–37.222	0.000	0.232	0.100
USD/CNY	–2.016	0.279	–1.742	0.410	1.206	0.010
log (USD/CNY)	–1.983	0.294	–1.718	0.422	1.193	0.010
R(USD/CNY)	–5.872	0.000	–37.532	0.000	0.234	0.100
USD/AUD	–2.280	0.179	–2.405	0.140	1.361	0.010
log (USD/AUD)	–2.305	0.170	–2.337	0.160	1.369	0.010
R(USD/AUD)	–35.700	0.000	–35.686	0.000	0.063	0.100
USD/INR	–1.058	0.732	–1.066	0.728	5.289	0.010
log (USD/INR)	–1.024	0.744	–1.039	0.739	5.292	0.010
R(USD/INR)	–36.835	0.000	–36.823	0.000	0.047	0.100

http://dx.doi.org/10.21511/imfi.22(3).2025.03

Variance analysis reveals that raw exchange rate series, especially for volatile pairs like USD/JPY and USD/INR, exhibit high variance that obscures patterns such as seasonality. Log transformations effectively reduce variance and address heteroscedasticity but do not isolate seasonality. To address this, CEEMDAN decomposition was applied, and the resulting IMFs (shown in Figures 6 to 10) were categorized into high-frequency components (IMFs 1-3) that capture short-term fluctuations likely driven by speculative activity, sentiment shifts, or microeconomic events. The medium-frequency components (IMFs 4-6 or 7, depending on the outcome) reflect cyclical behavior influenced by macroeconomic fundamentals such as interest rates and trade balances. The residual component represents the underlying long-term trend. Pairs like USD/EUR and USD/CNY exhibit smoother residuals and well-separated components, supporting the model's ability to distinguish between transient noise and structural dynamics. In contrast, more volatile pairs like USD/INR and USD/JPY reveal fewer stable trends and stronger high-frequency noise, aligning with their economic and geopolitical contexts. These decomposed features and the original time series data were fed into the TFT model. The model's attention mechanisms and variable selection dynamically assessed the relevance of these components, enhancing the forecasting process.

Table 4 compares the performance of the CEEMDAN-TFT model against conventional neural network models, including the Multilayer Perceptron (MLP), Recurrent Neural Network (RNN), Long Short-Term Memory (LSTM), and Convolutional Neural Network (CNN). The results clearly demonstrate that the TFT model outperforms these baseline models across a range of performance metrics, particularly showing the strongest predictive accuracy for the USD/EUR currency pair. This section provides a detailed analysis of both in-sample and out-of-sample results, underscoring the superior performance of the TFT model and its robustness in forecasting, with a notable emphasis on the USD/EUR pair.

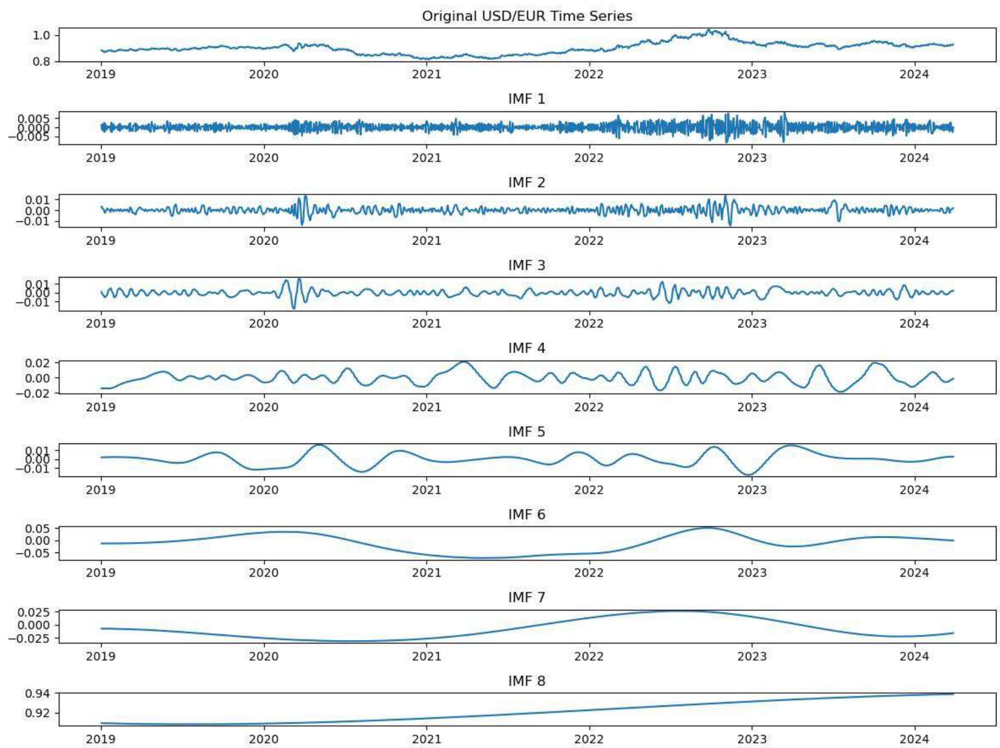

Note: IMFs 1–3 capture high-frequency noise and short-term volatility. IMFs 4–6 reflect medium-frequency market cycles, while the residual highlights the long-term trend component.

Figure 6. CEEMDAN decomposition of the USD/EUR exchange rate into Intrinsic Mode Functions (IMFs) and a residual trend

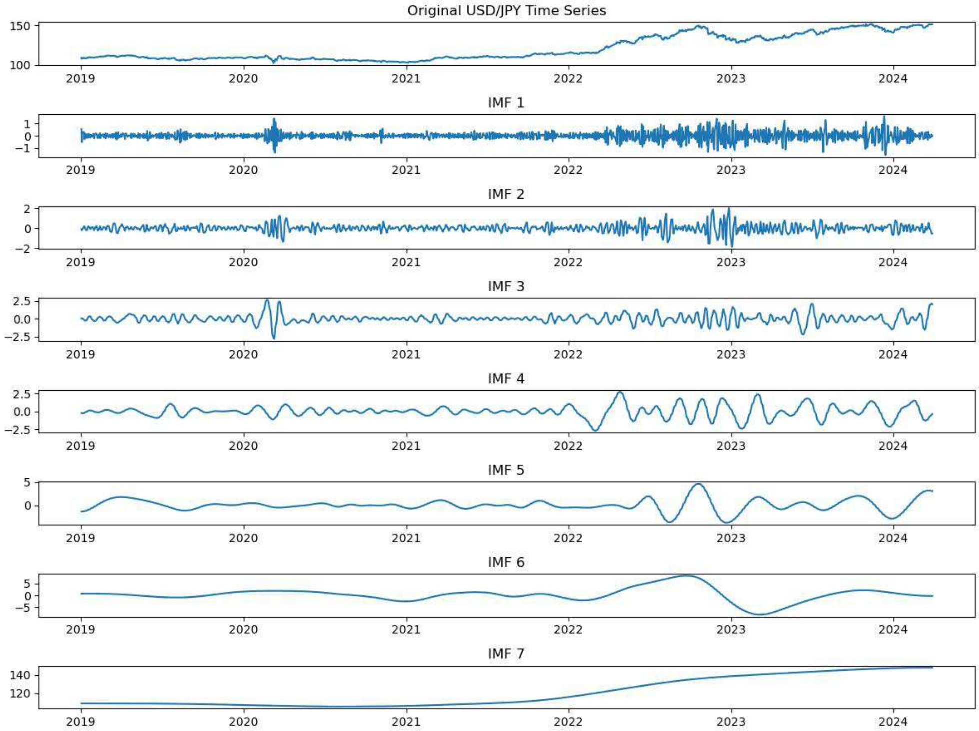

Note: The decomposition reveals pronounced medium-frequency oscillations and a relatively less stable long-term trend, consistent with the pair's higher volatility and sensitivity to macroeconomic events.

Figure 7. CEEMDAN decomposition of the USD/JPY exchange rate

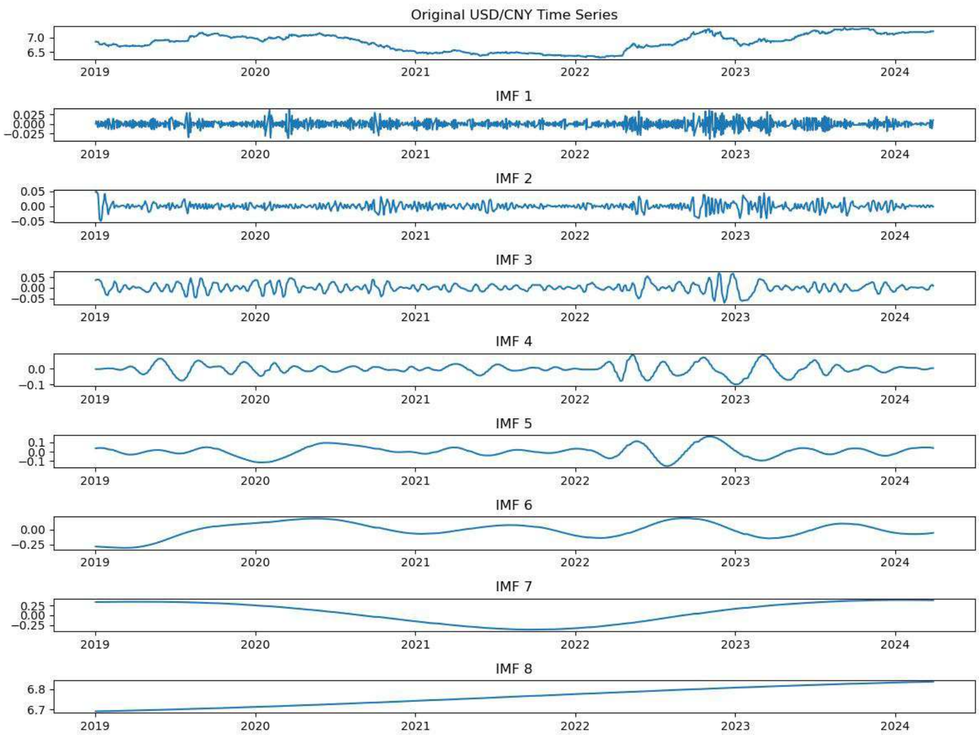

Note: The structure shows relatively dampened high-frequency IMFs, reflecting policy interventions and managed exchange rate behavior, with a smooth residual trend.

Figure 8. CEEMDAN decomposition of the USD/CNY exchange rate

http://dx.doi.org/10.21511/imfi.22(3).2025.03

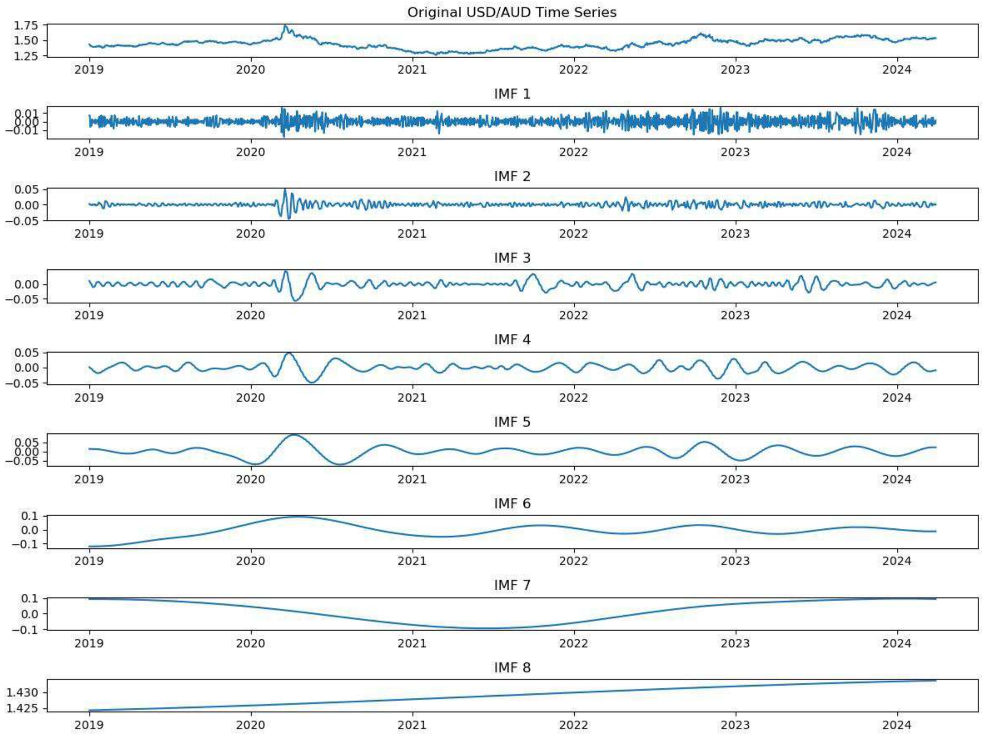

Note: The decomposition indicates strong medium-frequency fluctuations likely driven by commodity market cycles, with the residual trend aligning with global demand shifts.

Figure 9. CEEMDAN decomposition of the USD/AUD exchange rate

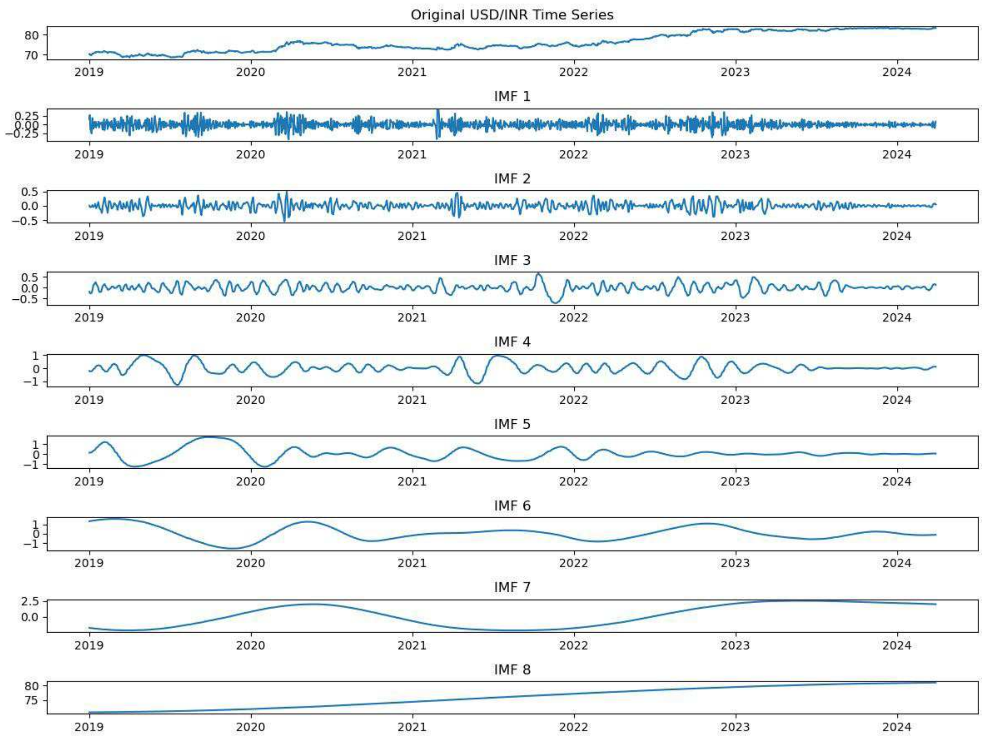

Note: High-frequency IMFs dominate the signal, indicative of frequent short-term volatility, while the residual trend captures a gradual appreciation trend with structural breaks.

Figure 10. CEEMDAN decomposition of the USD/INR exchange rate

To evaluate the performance of each model, a variety of error metrics, including Mean Squared Error (MSE), Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and R-squared (R^2). Additionally, robustness metrics such as Mean Absolute Percentage Error (MAPE), Symmetric Mean Absolute Percentage Error (SMAPE), and Mean Squared Logarithmic Error (MSLE). Further, the model performance was analyzed across different time horizons for a comprehensive assessment, specifically 1-day, 1-week, and 1-month RMSE. These metrics were chosen to provide a multifaceted view of each model's predictive accuracy and robustness.

The in-sample performance provides an initial understanding of how well each model fits the training data. The in-sample results across the five currency pairs (USD/EUR, USD/JPY, USD/CNY, USD/AUD, and USD/INR) are summarized below.

The USD/EUR pair consistently demonstrated the best model fit across all metrics. The CEEMDAN-TFT model achieved the lowest MSE (0.0104) and RMSE (0.1021) and the highest R^2 value (0.96), indicating that it explained 96% of the variance in the forex rates within the training data. These results were superior to those of the other models. The USD/JPY pair also showed strong performance with the TFT model, achieving an MSE (0.1046) and an RMSE (0.3234). The R^2 value (0.91) is slightly lower than USD/EUR but still indicates a strong relationship. Out-of-sample performance is critical for assessing a model's generalization ability to unseen data. Here, we analyze the results for the five currency pairs in the out-of-sample context shown in Table 5. The TFT model continued to outperform other models in the out-of-sample tests. It achieved an MSE of 0.0367 and RMSE of 0.1916, with an R^2 of 0.93, demonstrating minimal performance degradation compared to the in-sam-

Table 4. CEEMDAN-TFT and other model in-sample results

Fx-pair Model	In-Sample						
	MSE	MAE	RMSE	R^2	MAPE	SMAPE	MSLE
USD/EUR							
MLP	0.5745	0.5507	0.7580	0.8987	0.6420	0.9362	0.7517
RNN	0.7560	0.7560	0.8695	0.8662	0.7616	0.8511	0.1406
LSTM	0.6011	0.7081	0.7753	0.9199	0.6064	0.6695	0.7261
CNN	0.8324	0.5123	0.9124	0.8834	0.1015	0.6944	0.5163
TFT	0.0104	0.0225	0.1021	0.9612	0.0051	0.1701	0.0952
USD/JPY							
MLP	0.6119	0.9395	0.7822	0.8664	0.5487	0.6150	0.4390
RNN	0.4561	0.8852	0.6753	0.8142	0.6520	0.5401	0.2254
LSTM	0.5924	0.3465	0.7697	0.7705	0.7122	0.8774	0.8360
CNN	0.5651	0.9489	0.7517	0.8084	0.4254	0.6970	0.4090
TFT	0.1046	0.2977	0.3234	0.9102	0.2496	0.4595	0.0564
USD/CNY							
MLP	0.2220	0.4952	0.4712	0.8093	0.6576	0.8094	0.5492
RNN	0.5588	0.6625	0.7475	0.8201	0.3937	0.8671	0.6602
LSTM	0.5467	0.8849	0.7394	0.7751	0.4652	0.5113	0.9549
CNN	0.9395	0.8948	0.9693	0.8219	0.9730	0.7983	0.5544
TFT	0.3885	0.1960	0.6233	0.8253	0.2920	0.2020	0.4196
USD/AUD							
MLP	0.3887	0.2713	0.6234	0.7568	0.7948	0.8900	0.3560
RNN	0.7809	0.5427	0.8837	0.8022	0.5769	0.3756	0.5144
LSTM	0.7746	0.9869	0.8801	0.8987	0.1952	0.5783	0.0460
CNN	0.5055	0.8155	0.7110	0.7290	0.2808	0.4656	0.8555
TFT	0.2713	0.0740	0.5208	0.8159	0.1455	0.2865	0.4742
USD/INR							
MLP	0.8631	0.6233	0.9290	0.7636	0.9405	0.0305	0.4916
RNN	0.6110	0.3252	0.7817	0.6376	0.9149	0.8226	0.4732
LSTM	0.8872	0.4722	0.9419	0.7132	0.0155	0.6271	0.4985
CNN	0.7608	0.5613	0.8722	0.7938	0.4282	0.7700	0.6351
TFT	0.5227	0.2275	0.7230	0.8679	0.3636	0.4229	0.3746

http://dx.doi.org/10.21511/imfi.22(3).2025.03

ple results. All the above metrics were calculated using the test set, ensuring that the performance evaluation reflects generalization to unseen data.

Among the five currency pairs analyzed, the USD/EUR pair consistently exhibited the highest level of predictive accuracy across multiple error metrics, including MSE, RMSE, and R^2. The superior performance of the CEEMDAN-TFT model, for the USD/EUR pair, can be attributed to the fact that the USD/EUR exchange rate is widely regarded as the most liquid and extensively traded currency pair in the global forex market, characterized by high trading volume, greater data availability, and comparatively lower structural volatility (Fratzscher et al., 2019). These market characteristics can contribute to a more stable signal-to-noise ratio in the input data, which may facilitate better model learning and reduced prediction error. The model also achieved strong performance across other pairs, such as USD/JPY and USD/CNY, though with slightly higher error metrics, likely due to greater volatility, policy-driven interventions, and less transparent economic signals in those contexts. The result highlights the model's adaptability across different market conditions. Moreover, the out-of-sample RMSE for the USD/EUR pair of the model is lower than those reported in prior transformer-based forecasting studies (e.g., Lim et al., 2021; Cao et al., 2019), suggesting that the integration of CEEMDAN and sentiment variables may enhance prediction accuracy under certain market conditions. However, cross-study comparison should be interpreted with caution due to differences in data frequency, feature sets, and currency pairs.

Thereafter, the robustness and stability of the TFT model were analyzed by comparing the MAPE, SMAPE, and MSLE metrics. These metrics are particularly useful in assessing the model's performance across varying scales and under differ-

Table 5. CEEMDAN-TFT and other model out-of-sample results

Fx-pair Model	Out-of-Sample						
	MSE	MAE	RMSE	R^2	MAPE	SMAPE	MSLE
USD/EUR							
MLP	0.8314	0.6364	0.9118	0.9086	0.3841	0.3851	0.5314
RNN	0.2076	0.2493	0.4556	0.7556	0.8986	0.3169	0.6374
LSTM	0.2288	0.2770	0.4783	0.8612	0.2092	0.5568	0.9759
CNN	0.9297	0.8081	0.9642	0.8715	0.6635	0.6960	0.3230
TFT	0.0367	0.1866	0.1916	0.9393	0.1608	0.0972	0.2708
USD/JPY							
MLP	0.8074	0.8961	0.8986	0.8101	0.6919	0.9901	0.8785
RNN	0.9279	0.4271	0.9633	0.7607	0.3243	0.5183	0.9626
LSTM	0.2070	0.5107	0.4549	0.7221	0.2372	0.7408	0.6960
CNN	0.1199	0.3376	0.3462	0.8232	0.7465	0.7025	0.6733
TFT	0.1188	0.7030	0.3447	0.8972	0.2492	0.2936	0.3502
USD/CNY							
MLP	0.9624	0.2518	0.9810	0.7009	0.5683	0.8101	0.7146
RNN	0.5848	0.5369	0.7647	0.7803	0.3677	0.9132	0.2799
LSTM	0.3515	0.5786	0.5929	0.7396	0.2440	0.5015	0.7379
CNN	0.1449	0.4895	0.3807	0.7421	0.6931	0.6500	0.6117
TFT	0.1721	0.2616	0.4149	0.8282	0.1311	0.7958	0.3477
USD/AUD							
MLP	0.3678	0.6323	0.6065	0.7358	0.5026	0.3380	0.7578
RNN	0.8903	0.8353	0.9436	0.7865	0.4925	0.0940	0.1161
LSTM	0.4408	0.5909	0.6639	0.7166	0.7225	0.4359	0.1107
CNN	0.5121	0.2265	0.7156	0.7144	0.0243	0.5426	0.7037
TFT	0.2909	0.3867	0.5394	0.8038	0.1771	0.5908	0.0978
USD/INR							
MLP	0.3411	0.1135	0.5840	0.7773	0.9539	0.2373	0.4735
RNN	0.7579	0.6600	0.8706	0.6552	0.3702	0.3602	0.4339
LSTM	0.5297	0.2419	0.7278	0.6972	0.9283	0.5222	0.6159
CNN	0.9004	0.6331	0.9489	0.6492	0.9667	0.2158	0.2453
TFT	0.2260	0.2971	0.4753	0.8699	0.2530	0.1853	0.0259

ent error distributions. For the USD/EUR pair, the TFT model achieved MAPE: 0.0051, compared to 0.6420 (MLP), 0.7616 (RNN), 0.6695 (LSTM), and 0.1015 (CNN). SMAPE: 0.1701, compared to 0.9362 (MLP), 0.8511 (RNN), 0.6695 (LSTM), and 0.6944 (CNN). MSLE: 0.0952, compared to 0.7517 (MLP), 0.1406 (RNN), 0.7261 (LSTM), and 0.5163 (CNN). These results further validate the TFT model's ability to handle varying data distributions and minimize errors across different scales, making it a reliable choice for forex rate forecasting.

In addition to standard error metrics, we evaluated the models over different time horizons: 1-day, 1-week, and 1-month RMSE (shown in Figure 11). This analysis provides insight into the models' abilities to forecast over various temporal spans, which is crucial for practical trading and financial planning. The TFT model consistently showed the lowest RMSE across all time horizons for USD/EUR. These results underscore the TFT model's robustness across different forecasting horizons, making it particularly effective for both short-term and long-term forex predictions. The TFT model outperformed the baseline models across all time horizons for the remaining currency pairs. However, the margin of improvement was more pronounced in the short-term (1-day) and medium-term (1-week) forecasts. This indicates that while the TFT model is superior overall, its relative advantage is most evident in short to medium time frames. The forecast output is shown in Figures 12 to 16.

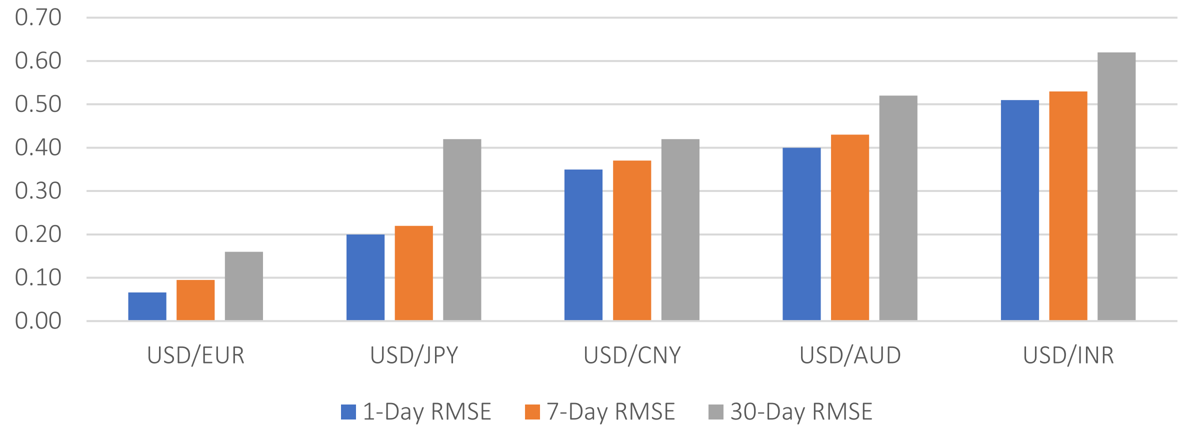

Figure 11. Model performance across time horizons

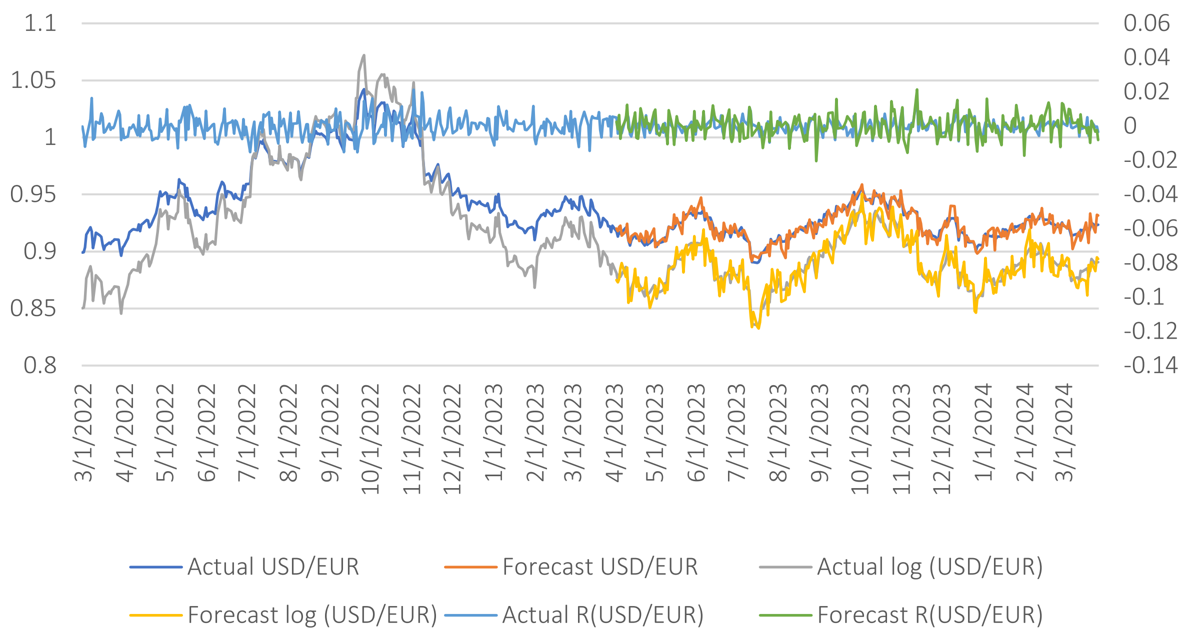

Figure 12. Model forecast for USD/EUR

http://dx.doi.org/10.21511/imfi.22(3).2025.03

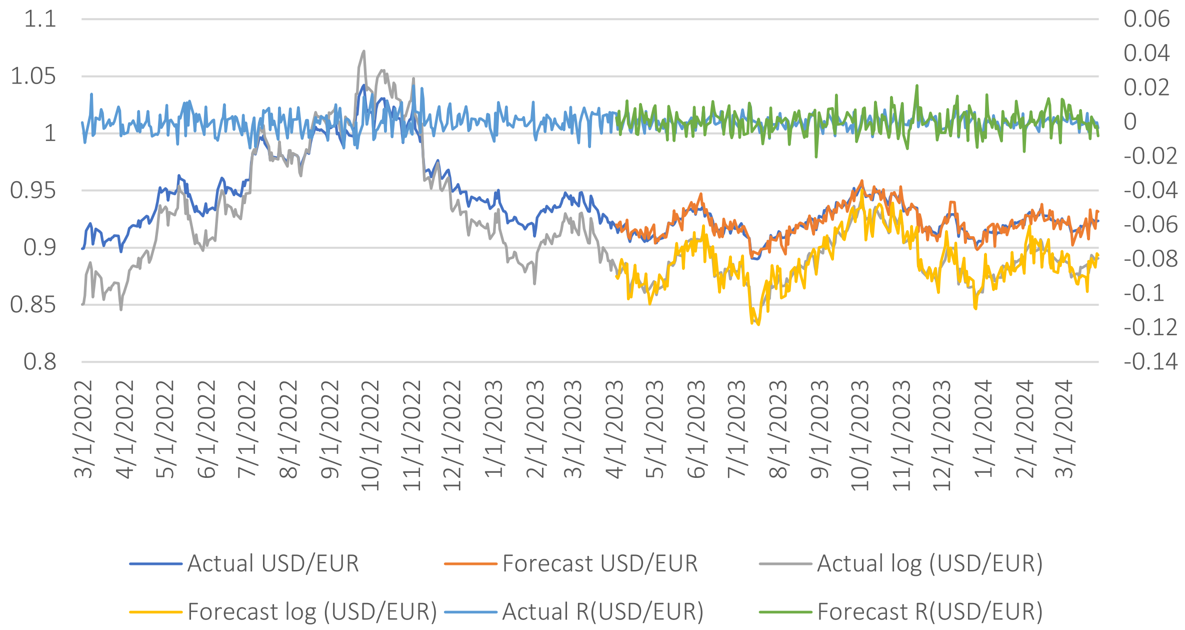

Figure 13. Model forecast for USD/JPY

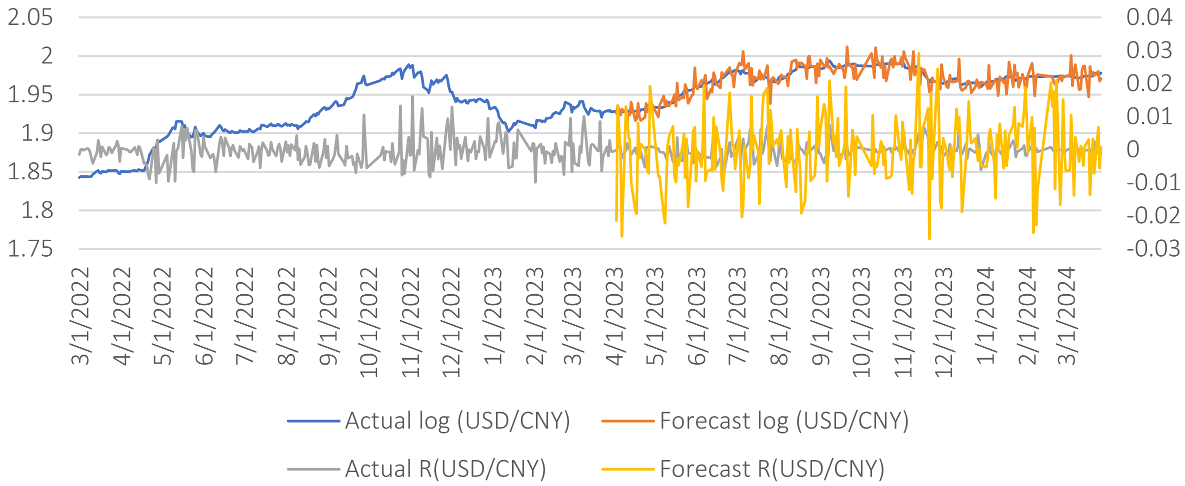

Figure 14. Model forecast for USD/CNY

The analysis revealed that the TFT model outperformed conventional neural network models across various performance metrics and provided valuable insights into the key variables influencing forex rate predictions. The model's ability to dynamically focus on the most impactful features at each time step provided deep insights into the driving forces behind currency movements, which varied across the different pairs (shown in Figure 17). For the USD/EUR currency pair, the TFT model identified the interest rates and investor sentiment scores as the most influential variables affecting the exchange rate. The model's attention mechanism emphasized these factors, particularly during periods of economic volatility and key financial announcements, underscoring their critical role in determining the strength of the EUR relative to the USD. This finding aligns with the observations of Heiden et al. (2013), who similarly noted that the EUR/USD market is especially sensitive to investor sentiment during specific periods. In the cases of USD/JPY and USD/CNY, while interest rates remained a significant factor, the model also identified GDP growth rate and trade balance as critical variables influencing the exchange rate. This outcome underscores the sensitivity of these

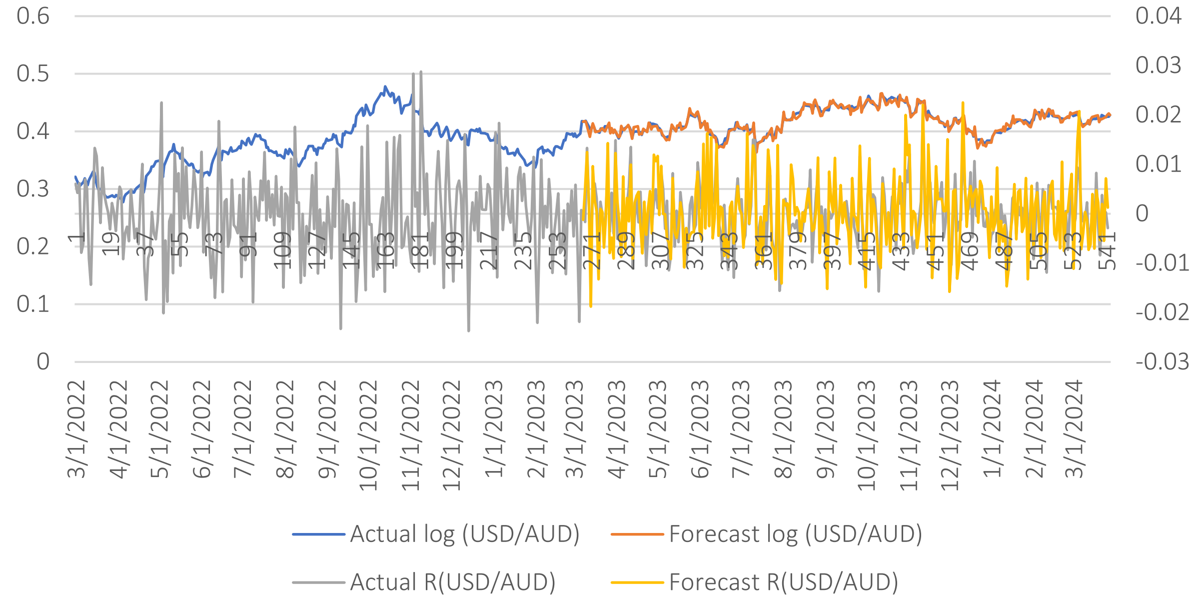

Figure 15. Model forecast for USD/AUD

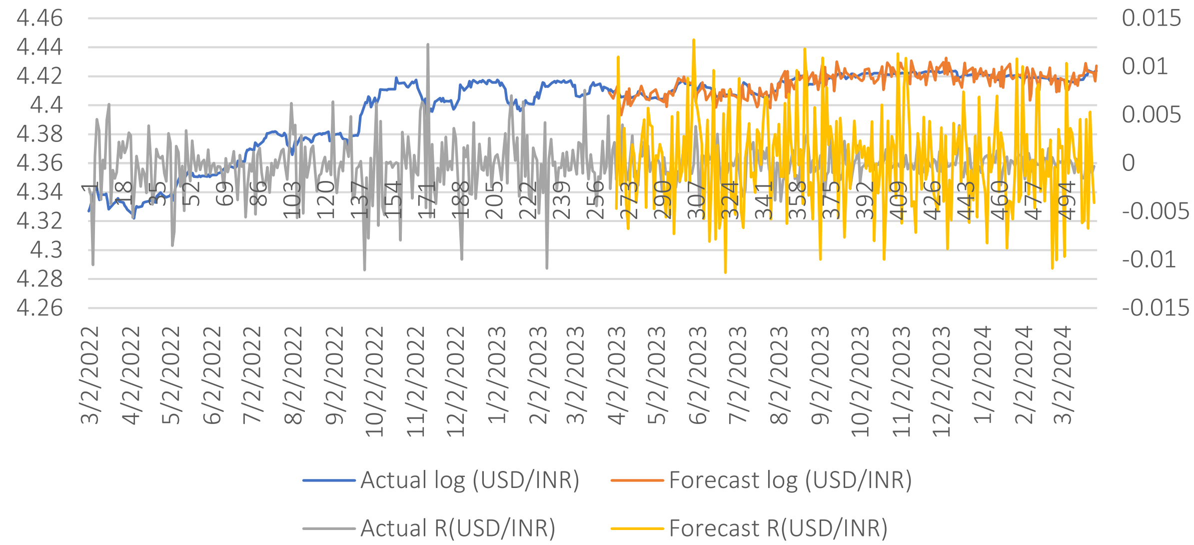

Figure 16. Model forecast for USD/INR

currencies to the economic performance of Japan and China, particularly given their export-driven economies. The prominence of these factors aligns with the existing literature, where Tang (2005) and Zhao et al. (2024) emphasized the impact of economic indicators on the valuation of these currencies. For the USD/AUD pair, the model identified commodity prices and investor sentiment as critical variables. This finding is intuitive given the Australian economy's heavy reliance on commodity exports, such as iron ore and gold, which makes the AUD particularly sensitive to global commodity price fluctuations (Simpson, 2002). The model's ability to weigh these factors accurately resulted in an in-sample R^2 of 0.81 and an out-of-sample R^2 of 0.80, with a 1-day RMSE of 0.4049. The USD/INR pair exhibited a slightly different dynamic, with the model highlighting inflation rates and investor sentiment as the key drivers of the exchange rate. India's status as an emerging market makes the INR particularly vulnerable to inflationary pressures, which can erode purchasing power and significantly impact foreign investment flows. This finding is consistent with the observations of Kumar & Aluvala (2020). This performance demonstrates the model's capability to adapt to the unique character-

http://dx.doi.org/10.21511/imfi.22(3).2025.03

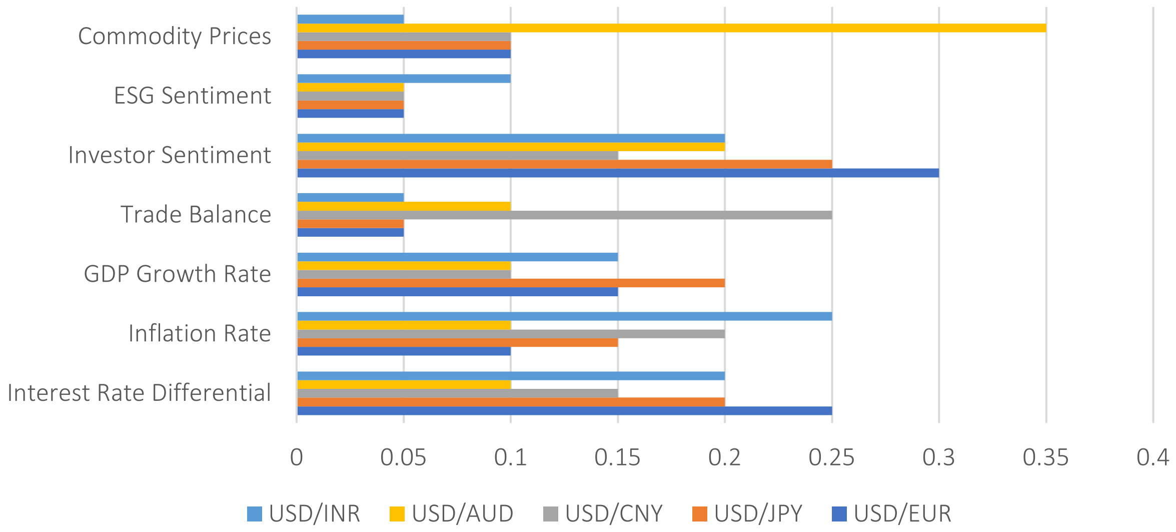

Figure 17. Attention weights of key variables across currency pairs

istics of emerging market currencies, providing accurate and context-specific forecasts for USD/INR, particularly in scenarios of economic instability or inflationary pressures.

The CEEMDAN-TFT model exhibited superior predictive accuracy across all five currency pairs, significantly outperforming traditional models such as MLP, RNN, LSTM, and CNN. The H1 hypothesis, which posited that the CEEMDAN-TFT model would outperform these neural network models, was validated through performance metrics including MSE, RMSE, and R^2. For example, the model achieved an R^2 of 0.9393 for the USD/EUR pair, compared to significantly lower values from the baseline models, confirming the hypothesis. Investor sentiment, analyzed using Twitter-derived data, was found to have a significant impact on forex rate predictions, especially during periods of heightened market volatility. This finding supports the acceptance of H2, as the inclusion of sentiment data notably improved predictive performance, with attention weights consistently highlighting sentiment variables during volatile episodes. For instance, during periods of central bank policy announcements, sentiment emerged as a dominant driver of predictions. Similarly, the integration of ESG sentiment scores into the model substantiated H3, demonstrating that ESG factors play a critical role in shaping currency movements. The attention mechanism of the TFT highlighted ESG variables as key drivers for pairs like USD/CNY and USD/AUD, particularly during sustainability-related news events. While the CEEMDAN decomposition technique enhanced data interpretability by isolating temporal patterns and reducing noise, the overall interpretability of the model remained less intuitive compared to simpler models like linear regression. This limitation led to the partial rejection of H4, which hypothesized that the integration of multi-scale decomposition with TFT would significantly enhance interpretability. Nonetheless, the multi-scale approach effectively captured critical trends and short-term fluctuations, contributing to improved forecasting accuracy.

CONCLUSION

This study explores the integration of multiscale time-series decomposition and attention-based deep learning to improve the forecasting of major forex currency pairs. By combining CEEMDAN and the Temporal Fusion Transformer (TFT), the proposed framework incorporates a wide range of economic indicators, investor sentiment, and ESG news sentiment to better capture the dynamic and nonlinear behavior of exchange rates. The empirical results suggest that the CEEMDAN-TFT model provides competitive predictive performance compared to conventional neural network models such as MLP, RNN, LSTM, and CNN. Across the five analyzed currency pairs, USD/EUR, USD/JPY, USD/CNY, USD/

AUD, and USD/INR, the model demonstrated adaptability to diverse economic contexts, with notably strong performance on the USD/EUR pair. This may be attributed to the liquidity and stability of the underlying data rather than inherent model superiority. The attention mechanism within TFT allowed for the dynamic identification of key drivers such as interest rate differentials, investor sentiment, and ESG factors, offering a degree of interpretability in line with current demands for explainable AI in financial modeling. While the results are promising, the study acknowledges several limitations. The interpretability of attention-based models remains less intuitive than traditional econometric methods, and the framework's effectiveness is partly contingent on the quality and granularity of the input data. Future research could extend this approach by exploring real-time applications, expanding the model to include more emerging market currencies, or comparing the performance across alternative transformer-based architectures.

AUTHOR CONTRIBUTIONS

Conceptualization: Sougata Banerjee.
Data curation: Sougata Banerjee.
Formal analysis: Sougata Banerjee.
Investigation: Sougata Banerjee.
Methodology: Sougata Banerjee.
Project administration: Sougata Banerjee.
Resources: Sougata Banerjee.
Software: Sougata Banerjee.
Supervision: Sougata Banerjee.
Validation: Sougata Banerjee.
Visualization: Sougata Banerjee.
Writing – original draft: Sougata Banerjee.
Writing – review & editing: Sougata Banerjee.

REFERENCES

1. Abdel Magid, A., Hussainey, K., De Andrés, J., & Lorca, P. (2023). The moderating role of online social media in the relationship between corporate social responsibility disclosure and investment decisions: evidence from Egypt. *International Journal of Financial Studies, 11*(2), 60. https://doi.org/10.3390/ijfs11020060

2. Abraham, B. (1982). Temporal Aggregation and Time Series. *International Statistical Review / Revue Internationale de Statistique, 50*(3), 285. https://doi.org/10.2307/1402498

3. Alexander, D., & Thomas, L. R. (1987). Monetary/asset models of exchange rate determination. How well have they performed in the 1980's? *International Journal of Forecasting, 3*(1), 53-64. https://doi.org/10.1016/0169-2070(87)90078-1

4. Bekiros, S., & Marcellino, M. (2013). The multiscale causal dynamics of foreign exchange markets. *Journal of International Money and Finance, 33,* 282-305. https://doi.org/10.1016/j.jimonfin.2012.11.016

5. Bhattacharya, B. (2012). *Causal relationship between stock market and exchange rate, foreign exchange reserves and value of trade balance in India: An Empirical analysis.* Retrieved from http://oii.igidr.ac.in:8080/jspui/bitstream/2275/185/1/basabi.pdf

6. Bose, D., Aithal, P., & Roy, S. (2021). Survey of Twitter Viewpoint on Application of Drugs by VADER Sentiment Analysis among Distinct Countries. *International Journal of Management*. Retrieved from https://srinivaspublication.com/wp-content/uploads/2021/03/6.-Survey_Fullpaper.pdf

7. Cao, J., Li, Z., & Li, J. (2019). Financial time series forecasting model based on CEEMDAN and LSTM. *Physica A: Statistical Mechanics and Its Applications, 519,* 127-139. https://doi.org/10.1016/j.physa.2018.11.061

8. Chen, Y. C., Rogoff, K. S., & Rossi, B. (2010). Can exchange rates forecast commodity prices? *Quarterly Journal of Economics, 125*(3), 1145-1194. Retrieved from https://www.nber.org/system/files/working_papers/w13901/w13901.pdf

9. Chen, Y., Zheng, Y., Lv, G., Gu, Y., & Zhang, W. (2024). Quantifying the effect of ESG-related news on Chinese stock movements. *Journal of Environmental Management, 354,* 120301. https://doi.org/10.1016/j.jenvman.2024.120301

10. De Myttenaere, A., Golden, B., Le Grand, B., & Rossi, F. (2016).

http://dx.doi.org/10.21511/imfi.22(3).2025.03

Mean absolute percentage error for regression models. *Neurocomputing, 192,* 38-48. https://doi.org/10.1016/j.neucom.2015.12.114

11. Embrechts, M. J. (1995). Forecasting foreign exchange rates with artificial neural networks. *Intelligent Engineering Systems Through Artificial Neural Networks, 5*(1), 771-778.

12. Engle, R. (2001). GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. *Journal of Economic Perspectives, 15*(4), 157-168. Retrieved from https://ideas.repec.org/a/aea/jecper/v15y2001i4p157-168.html

13. Fama, E. F. (1970). Efficient Capital Markets. *The Journal of Finance, 25*(2), 383-417. https://doi.org/10.2307/2325486

14. Fratzscher, M., Gloede, O., Menkhoff, L., Sarno, L., & Stöhr, T. (2019). When is Foreign exchange intervention effective? Evidence from 33 countries. *American Economic Journal: Macroeconomics, 11*(1), 132-156. Retrieved from https://www.aeaweb.org/articles?id=10.1257/mac.20150317

15. Guo, W., & Chen, Z. (2023). China–US economic and trade relations, trade news, and short-term fluctuation of the RMB exchange rate. *Review of International Economics, 31*(1), 180-203. https://ideas.repec.org/a/bla/reviec/v31y2023i1p180-203.html

16. Heiden, S., Klein, C., & Zwergel, B. (2013). Beyond fundamentals: Investor sentiment and exchange rate forecasting. *European Financial Management, 19*(3), 558-578. Retrieved from https://ideas.repec.org/a/bla/eufman/v19y2013i3p558-578.html

17. Hong, C., & Zu, W. (2024). A Twofold Model for Exchange Rate Forecasting: Combining Fundamentals and Market Dynamics. *Journal of Business and Economic Options, 7*(4), 42-52. Retrieved from https://resdojournals.com/index.php/jbeo/article/view/393

18. Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Networks, 4*(2), 251-257. https://doi.org/10.1016/0893-6080(91)90009-T

19. Huang, X., Wang, Y., & Li, H. (2024). Exploring the asymmetric influence of economic policy uncertainty on the nonlinear relationship between exchange rate and carbon prices in China. *North American Journal of Economics and Finance, 73*. Retrieved from https://ideas.repec.org/a/eee/ecofin/v73y2024ics1062940824000913.html

20. Huy, P. C., Minh, N. Q., Tien, N. D., & Anh, T. T. Q. (2022). Short-Term Electricity Load Forecasting Based on Temporal Fusion Transformer Model. *IEEE Access, 10,* 106296-106304.

21. Ito, T., & Sato, K. (2008). Exchange rate changes and inflation in post-crisis Asian economies: Vector autoregression analysis of the exchange rate pass-through. *Journal of Money, Credit and Banking, 40*(7), 1407-1438. https://doi.org/10.1111/j.1538-4616.2008.00165.x

22. Jan, J. H., & Gopalaswamy, A. K. (2019). Identifying factors in currency exchange rate estimation: a study on AUD against USD. *Journal of Advances in Management Research, 16*(4), 436-452. Retrieved from https://ideas.repec.org/a/eme/jamrpp/jamr-09-2018-0084.html

23. Kourentzes, N., Petropoulos, F., & Trapero, J. R. (2014). Improving forecasting by estimating time series structural components across multiple frequencies. *International Journal of Forecasting, 30*(2), 291-302. https://doi.org/10.1016/j.ijforecast.2013.09.006

24. Kumar, J. S., & Aluvala, R. (2020). Impact of selected economic variables on the USD/INR exchange rate. *International Journal of Management (IJM), 11*(7). https://iaeme.com/MasterAdmin/Journal_uploads/IJM/VOLUME_11_ISSUE_7/IJM_11_07_064.pdf

25. Lim, B., Arık, S., Loeff, N., & Pfister, T. (2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting, 37*(4), 1748-1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

26. MacDonald, R. (2007). Exchange rate economics: Theories and evidence. In *Exchange Rate Economics: Theories and Evidence*. Routledge Taylor & Francis Group. Retrieved from https://www.taylorfrancis.com/books/mono/10.4324/9780203380185/exchange-rate-economics-ronald-macdonald

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

28. Mehta, P., Pandya, S., & Kotecha, K. (2021). Harvesting social media sentiment analysis to enhance stock market prediction using deep learning. *PeerJ Computer Science, 7,* 1-21. Retrieved from https://peerj.com/articles/cs-476/

29. Menkhoff, L., & Taylor, M. P. (2007). The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature, 45*(4), 936-972. Retrieved from https://ideas.repec.org/p/wrk/warwec/769.html

30. Mohammadipour, M., & Boylan, J. E. (2012). Forecast horizon aggregation in integer autoregressive moving average (INARMA) models. *Omega, 40*(6), 703-712. https://doi.org/10.1016/j.omega.2011.08.008

31. Panda, M. M., Panda, S. N., & Pattnaik, P. K. (2022). Forecasting foreign currency exchange rate using convolutional neural network. *International Journal of Advanced Computer Science and Applications, 13*(2), 607-616. Retrieved from https://thesai.org/Downloads/Volume13No2/Paper_72-Forecasting_Foreign_Currency_Exchange_Rate.pdf

32. Plakandaras, V., Papadimitriou, T., & Gogas, P. (2015). Forecasting daily and monthly exchange rates with machine learning techniques. *Journal of Forecasting, 34*(7), 560-573. https://doi.org/10.1002/for.2354

33. Safari Bideskan, S., Mehrazeen, A. R., Shorvarzi, M., & Masihabadi, A. (2022). Have hashtags and cashtags caused a slight reaction to stock returns in financial statements? Has the information content of the financial statements been lost? Case study: S & P500 companies. *International Journal of Nonlinear Analysis and Applications, 13*(1), 3029-3057. Retrieved from https://ijnaa.semnan.ac.ir/article_6041.html

34. Sharma, S., Aggarwal, V., & Yadav, M. P. (2021). Comparison of linear and non-linear GARCH models for forecasting volatility of select emerging countries. *Journal of Advances in Management Research, 18*(4), 526-547. Retrieved from https://ideas.repec.org/a/eme/jamrpp/jamr-07-2020-0152.html

35. Simpson, J. L. (2002). The relationship between commodity prices and the Australian dollar. *SSRN Electronic Journal*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=314872

36. Steuer, S., & Tröger, T. H. (2022). The role of disclosure in green finance. *Journal of Financial Regulation, 8*(1), 1-50. Retrieved from https://academic.oup.com/jfr/article/8/1/1/6539733

37. Suman, S., & Raju, G. (2024). Dynamic Interaction Between the Nifty Index and Exchange Rate: Evidence from VAR Analysis. *Educational Administration: Theory and Practice, 30*(5), 5479-5487. Retrieved from https://kuey.net/index.php/kuey/article/view/3805

38. Tang, T. C. (2005). Does exchange rate volatility matter for the balancing item of balance of payments accounts in Japan? An empirical note. *International Review of Economics and Business, 52*(4), 581-590. Retrieved from https://research.monash.edu/en/publications/does-exchange-rate-volatility-matter-for-the-balancing-item-of-ba

39. Torres, M. E., Colominas, M. A., Schlotthauer, G., & Flandrin, P. (2011). A complete ensemble empirical mode decomposition with adaptive noise. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP),* 4144-4147. Retrieved from https://ieeexplore.ieee.org/abstract/document/5947265

40. Weigend, A. S., Huberman, B. A., & Rumelhart, D. E. (1992). Predicting Sunspots and Exchange Rates with Connectionist Networks. In *Nonlinear modeling and forecasting, 12*(1991), 395-432). Retrieved from https://www.researchgate.net/publication/313334261_Predicting_sunspots_and_exchange_rates_with_connectionist_networks_nonlinear_modeling_and_forecasting

41. Yilmaz, F. M., & Arabaci, O. (2021). Should deep learning models be in high demand, or should they simply be a very hot topic? A comprehensive study for exchange rate forecasting. *Computational Economics, 57*(1), 217-245. Retrieved from https://ideas.repec.org/a/kap/compec/v57y2021i1d10.1007_s10614-020-10047-9.html

42. Zhang, H., Zou, Y., Yang, X., & Yang, H. (2022). A temporal fusion transformer for short-term freeway traffic speed multistep prediction. *Neurocomputing, 500,* 329-340. https://doi.org/10.1016/j.neucom.2022.05.083

43. Zhang, W., Qu, Z., Zhang, K., Mao, W., Ma, Y., & Fan, X. (2017). A combined model based on CEEMDAN and modified flower pollination algorithm for wind speed forecasting. *Energy Conversion and Management, 136,* 439-451. https://doi.org/10.1016/j.enconman.2017.01.022

44. Zhao, L., Yao, D., Yang, B., & Zhu, K. (2024). The Impact of China-United States Trade War on the Relationship Among the Exchange Rates: The Case of China, Korea, and Japan. *Machine Learning for Econometrics and Related Topics,* 489-499. Retrieved from https://link.springer.com/chapter/10.1007/978-3-031-43601-7_32

http://dx.doi.org/10.21511/imfi.22(3).2025.03