A CEEMDAN-LSTM MODEL FOR FORECASTING THE USD/DZD EXCHANGE RATE[1]

1 Abdelkader Sahed, University Centre of Maghnia, Department of Economics, Algeria
2 Mohammed Mekidiche, University Centre of Maghnia, Department of Economics, Algeria
3 Hacen Kahoui, University Centre of Maghnia, Department of Economics, Algeria
*Corresponding author´s e-mail: a.sahed@cu-maghnia.dz
1 ORCID ID: 0009-0003-2776-0448
2 ORCID ID: 0000-0001-5001-9250
3 ORCID ID: 0000-0003-1760-2940

ARTICLE INFO
Original Scientific Paper
Received: 18.12.2025.
Revised: 21.05.2026.
Accepted: 03.06.2026.
doi:10.63356/ace.2026.003
UDK
339.13.017(73):336.748.3(65)
COBISS.RS-ID 144551681

Keywords: *CEEMDAN, LSTM, exchange rate forecasting, financial time series, deep learning, risk management*

JEL Classification: C45, C53, F31, G17

ABSTRACT

This study develops a hybrid forecasting model for the USD/DZD exchange rate by combining Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) and Long Short-Term Memory (LSTM) networks to address high volatility and complex temporal dependencies in currency markets. Using 310 monthly observations, the CEEMDAN procedure decomposes the series into five frequency components and a residual, which are modeled by component-specific LSTM networks. The proposed CEEMDAN-LSTM model achieved the lowest forecast errors among the tested models, with a MAPE of 0.4782%, outperforming traditional LSTM and SVM benchmarks. The 12-month forecast suggests relative exchange rate stability, with a slight decline of about 0.48%. The results indicate that decomposing the original series before LSTM modeling improves predictive accuracy by separating short-term noise from medium- and long-term dynamics. The proposed framework may support exchange-rate risk management, hedging decisions, and short- to medium-term planning in emerging-market settings.

1. INTRODUCTION

Exchange rates represent one of the fundamental components of the macroeconomic structure, given their direct impact on foreign trade patterns, capital flows, balance of payments, and a country's international competitiveness (Plakandaras et al., 2015). The USD/DZD exchange rate stands out as a strategic indicator in the Algerian context, due to the dollar's dominance in pricing key exports, particularly hydrocarbons, on one hand, and the significant dependence of commodity and service imports on it, on the other. In this context, the ability to accurately forecast exchange rate movements becomes a critical tool for monetary authorities, policymakers, and financial market operators—not only for managing and hedging against exchange rate volatility risks, but also for designing more targeted monetary policies and enhancing the efficiency of short- and medium-term financial and budgetary planning, which ultimately reflects on macroeconomic stability.

However, exchange rate time series, particularly in resource-dependent economies, typically exhibit high degrees of instability, volatility, non-linearity, and complex interactions with various financial, economic, and geopolitical factors. These characteristics render many traditional time series analysis models, such as linear regression or ARIMA, limited in their ability to represent the dynamic relationships governing exchange rate movements (De Gooijer & Hyndman, 2006). This has encouraged the use of more flexible methods, including Empirical Mode Decomposition (EMD) and Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN). EMD, originally proposed by Huang et al. (1998), decomposes non-stationary signals into intrinsic mode functions (IMFs) through an iterative process. CEEMDAN, introduced by Torres et al. (2011), was developed to reduce mode mixing and generate more stable components, thereby improving subsequent forecasting performance (Cao et al., 2019; Wu & Huang, 2009).

Parallel to this development, recent years have witnessed significant expansion in the application of artificial intelligence techniques, particularly machine learning and deep learning models, in financial data and exchange rate forecasting. Support Vector Machines (SVM), pioneered by Vapnik (1995), have proven effective in modeling non-linear relationships within exchange rate data through kernel methods and margin maximization principles. Long Short-Term Memory (LSTM) neural networks, introduced by Hochreiter & Schmidhuber (1997), have demonstrated particular effectiveness in capturing long-term temporal dependencies within series, achieving superior results compared to many traditional models in financial market and exchange rate forecasting (Zhang, 2018). The LSTM architecture addresses the vanishing gradient problem inherent in recurrent neural networks, enabling the network to learn and retain information over extended temporal sequences (Hochreiter & Schmidhuber, 1997). The convergence of these two paths—time series decomposition methods like CEEMDAN and deep learning models like LSTM—has led to the emergence of hybrid models that leverage the strengths of both approaches. In such hybrid frameworks, mode decomposition reduces noise and instability in the original signal, while deep networks extract complex temporal patterns from the resulting components (Fan et al., 2021).

Building on this theoretical and methodological background, this study addresses the following main question: How can a hybrid model that integrates CEEMDAN with LSTM networks improve the accuracy of USD/DZD exchange rate forecasting relative to benchmark single-model approaches such as SVM and LSTM? This problem falls within a broader effort to test the viability of hybrid models in an economic environment with distinctive characteristics such as the Algerian economy, where external and internal factors jointly shape exchange rate movements and make prediction particularly challenging when relying on a single modeling technique.

In light of this, the study aims to achieve several interconnected objectives. First, it decomposes the USD/DZD exchange rate series using CEEMDAN into more stable and modelable sub-components. Second, it constructs benchmark SVM and LSTM forecasting models and evaluates them using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE). Third, it proposes a CEEMDAN-LSTM hybrid model that feeds the decomposed components into LSTM networks, combining CEEMDAN's capacity to reduce noise and instability with LSTM's ability to capture complex temporal dependencies. Fourth, it compares the hybrid model with the benchmark models in order to assess the forecasting gains generated by the proposed methodology (Fan et al., 2021).

With this framework, the study seeks to contribute to filling part of the research gap in the literature on exchange rate forecasting in emerging economies by presenting an applied model that combines advanced time series analysis techniques with artificial intelligence tools. This could provide a practical framework beneficial to policymakers, monetary authorities, and financial institutions in enhancing the efficiency of exchange rate volatility risk management in the Algerian context, with the potential to generalise the methodology to other currencies and markets with similar characteristics. Furthermore, by demonstrating the superior performance of hybrid approaches over standalone models, this research contributes to the growing body of evidence supporting the use of ensemble and hybrid methods in financial forecasting (Ren et al., 2016).

To systematically address the research question and achieve the stated objectives, this paper is organised as follows. The second section reviews the literature on exchange rate forecasting, covering both traditional econometric approaches and contemporary machine-learning techniques. The third section presents the data, preprocessing steps, and the architecture of the proposed CEEMDAN-LSTM model. The fourth section presents and discusses the empirical results, comparing the forecasting performance of the hybrid model with benchmark LSTM and SVM models using MAE, RMSE, and MAPE. The fifth section concludes and outlines the main limitations and directions for future research.

2. LITERATURE REVIEW

Forecasting foreign exchange rates represents one of the primary challenges in economics and finance, as exchange rates are influenced by multiple factors including inflation, interest rates, monetary policies, and geopolitical events. Accurate forecasting helps in making informed investment decisions, managing financial risks, and promoting economic stability. In this section, we review previous studies related to exchange rate forecasting models, focusing on traditional models, artificial intelligence models, and hybrid models that combine signal decomposition techniques (such as EMD or CEEMDAN) with neural networks like LSTM. We emphasise studies addressing different currencies, while highlighting the gap in applying the CEEMDAN-LSTM model to the US Dollar against the Algerian Dinar (USD/DZD) exchange rate.

Many previous studies rely on traditional time series models, such as ARIMA and SARIMA, which focus on linear and seasonal patterns in the data. For instance, Al-Gounmeein and Ismail (2020) employed a seasonal ARIMA model to forecast the Jordanian dinar against the US dollar using monthly data. Their results showed that SARIMA (1,0,1)(1,0,0)12 outperformed ARIMA (1,0,1) in terms of MAPE, RMSE, and MAE, while also highlighting the limitations of linear models in capturing nonlinear exchange-rate fluctuations.

Similarly, Darvas & Schepp (2025) proposed a monetary model based on the rational-expectations present-value framework for forecasting the pound sterling against the US dollar. Although their model outperformed the random walk benchmark and generated positive economic and statistical returns, the authors also emphasised the persistent difficulty of forecasting major exchange rates under unexpected volatility.

In another context, Flores-Sosa et al. (2022) used a multiple linear regression-heavy ordered weighted average (MLR-HOWA) framework to forecast five Latin American currencies against the US dollar. Their model reduced forecasting error relative to conventional linear regression and highlighted the value of multi-scenario analysis.

With the advancement of artificial intelligence techniques, studies have increasingly adopted neural networks such as ANN, LSTM, and BiLSTM to handle the nonlinear and non-stationary nature of exchange rates. Kamouh (2026) used artificial neural networks to forecast the Egyptian pound against the US dollar by relying on variables such as external debt, GDP, and inflation, and reported higher accuracy than traditional models. García et al. (2024) compared LSTM and BiLSTM models for exchange rates of major currencies against the US dollar, including Bitcoin futures contracts, and found BiLSTM particularly effective for short-term forecasting.

Amri et al. (2025) proposed a deep neural network with a multi-output sliding-window approach for forecasting the Indonesian rupiah against the US dollar and reported improved accuracy through better handling of multiple temporal patterns. Liu et al. (2023) developed a hybrid CNN-STLSTM-AM model for USD/RMB forecasting and found that it outperformed CNN, SVR, LSTM, and GRU-LSTM benchmarks.

To address noise in time series, some studies have adopted decomposition techniques such as Empirical Mode Decomposition (EMD) or CEEMDAN combined with artificial intelligence models. Lin et al. (2012) used an EMD-LSSVR model to forecast multiple currency exchange rates. By decomposing the original series into IMFs and residuals before prediction, the model outperformed both EMD-ARIMA and standalone LSSVR.

Adesina and Obokoh (2025) proposed a SARIMA-LSTM framework for forecasting the South African rand against the dollar, euro, and yuan, reporting gains over ARIMA, SARIMA, GRU, RNN, and standalone LSTM specifications. Lyócsa et al. (2024) examined day-ahead expected shortfall for the EUR/USD rate and showed that implied volatility can improve forecasting accuracy, especially during turbulent periods.

Despite progress in hybrid models, most previous studies have focused on major currencies such as USD/RMB, USD/GBP, or USD/EGP, with few studies on emerging currencies like the Algerian Dinar (DZD). Additionally, the use of CEEMDAN (an improvement over EMD for noise reduction) with LSTM has not been widely applied to USD/DZD. For example, none of the previous studies specifically addressed the USD/DZD exchange rate, which is influenced by local factors such as oil prices and monetary policies in Algeria. This research fills this gap by proposing a CEEMDAN-LSTM model, which combines complete signal decomposition (CEEMDAN) to extract principal components and LSTM to handle long-term temporal dependencies, thereby improving accuracy in forecasting the USD/DZD exchange rate compared to previous models.

3. MATERIALS AND METHODS

3.1 EMD Decomposition

Empirical Mode Decomposition (EMD) is an adaptive time series analysis technique based on the Hilbert-Huang Transform (HHT), specifically designed to handle nonlinear and non-stationary temporal data (Huang et al., 1998). The fundamental principle of EMD lies in decomposing a time series into a set of oscillatory functions called Intrinsic Mode Functions (IMFs), where each function must satisfy two essential conditions: first, the number of extrema (sum of maxima and minima) and the number of zero-crossings must be equal or differ by at most one throughout the entire series; and second, the mean value of the envelopes defined by the local maxima and minima must equal zero at all points. These requirements ensure the extraction of meaningful intrinsic mode functions; otherwise, blind application of the technique to any data may result in physically meaningless harmonics (Huang et al., 1999).

The decomposition process is performed through a systematic procedure called the "Sifting Procedure" applied to any data series $x(t)$, where the process begins by identifying all local extrema of the series, then these values are connected using a curve to create upper and lower envelopes (Yu et al., 2008). Subsequently, the point-wise mean of the envelopes is calculated as:

$$m(t) = \frac{x_{up}(t) + x_{low}(t)}{2}$$

and the details are extracted as:

$$c(t) = x(t) - m(t)$$

This process is repeated until a stopping criterion is satisfied, defined by three conditions related to predetermined thresholds (Huang et al., 1998). The extraction of intrinsic mode functions continues by applying the sifting procedure to the remaining residue until the final residue $r_n(t)$ becomes a monotonic function or contains no more than one local extremum. At the end of the procedure, the original series $x(t)$ can be mathematically expressed as the sum of all extracted intrinsic mode functions and the final residue:

$$x(t) = \sum_{j} c_j(t) + r_n(t)$$

where $c_j(t)$ represents the intrinsic mode functions that are nearly orthogonal to each other and have means approaching zero, while $r_n(t)$ represents the main trend of the time series. The EMD technique has several clear advantages compared to traditional decomposition methods such as Fourier analysis and Wavelet analysis. First, it is easy to understand and implement; second, it automatically and adaptively selects oscillations within the time series from the series itself, making it powerful in analysing nonlinear and non-stationary time series; and third, it follows a data-driven approach where the data speaks for itself without imposing prior assumptions. The most important advantage of EMD is that it does not require specifying a basis filter function before the decomposition process, unlike Wavelet analysis which requires pre-specifying this function—a difficult task for unknown series (Li, 2006). Studies have also demonstrated that EMD performs better in describing instantaneous frequencies on the local time scale compared to Wavelet analysis and Fourier analysis (Huang et al., 1999), as Fourier analysis when applied to nonlinear time series often produces large sets of physically meaningless harmonics as the degree of nonlinearity and non-stationarity in the series increases. The sifting procedure can be easily implemented using MATLAB software (Yu et al., 2008), making it accessible for practical applications in exchange rate forecasting and other time series analysis tasks.

3.2 Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN)

The Complete Ensemble Empirical Mode Decomposition with Adaptive Noise technique was developed by Torres et al. (2011). CEEMDAN represents a significant breakthrough in analysing nonlinear and non-stationary time series signals, with applications including building energy consumption data analysis.

CEEMDAN was specifically designed to address limitations found in previous decomposition methods, particularly Empirical Mode Decomposition (EMD) and Ensemble Empirical Mode Decomposition (EEMD). While EMD encounters mode mixing problems caused by intermittent signals, EEMD attempts to resolve this by incorporating white noise but introduces residual noise contamination, especially with low ensemble numbers, and requires increased computational time with higher ensemble sizes. CEEMDAN overcomes these challenges by incorporating adaptive noise at each decomposition stage, thereby improving accuracy while reducing computational complexity. The core objective of CEEMDAN is to decompose an original signal into relatively stable components called Intrinsic Mode Functions (IMFs), plus a residual component. This decomposition facilitates individual prediction of each component, making it particularly valuable for forecasting applications.

The method begins by defining the following parameters:

- $x[n]$: the original signal
- w^i: a white noise series following a normal distribution $N[0,1]$
- ε_k: noise amplitude coefficients
- $Ej(\cdot)$: a function that extracts the j-th mode using EMD.

The decomposition process involves several systematic steps that progressively extract IMFs from the signal:

Step 1: Initial Noise Addition

White noise is added to the original signal:

$$x[n]+\varepsilon_0 w^i[n]$$

Step 2: First IMF Extraction

I realisations of the noise-added signal are decomposed using EMD to extract the first modes. The first IMF is calculated as the average:

$$IMF_1[n]=\frac{1}{I}\sum_{i=1}^{I} IMF_1^i[n]=\overline{IMF_1}[n]$$

Step 3: First Residue Calculation

$$r_1[n]=x[n]-IMF_1[n]$$

Step 4: Second IMF Extraction

The process continues by decomposing $r_1[n]+\varepsilon_1 E_1(w^i[n])$ for $i = 1,\ldots,I$ using EMD to extract the first mode, then computing the second IMF:

$$IMF_2[n]=\frac{1}{I}\sum_{i=1}^{I} E_1(r_1[n]+\varepsilon_1 E_1(w^i[n]))$$

Step 5: Iterative Residue Update

For each $k = 2,...,K$, the k-th residue is calculated:

$$r_k[n]=r_{k-1}[n]-IMF_k[n]$$

Step 6: Subsequent IMF Extraction

The decomposition continues by processing $r_k[n]+\varepsilon_k E_k(w^i[n])$ for $i = 1,\ldots,I$, computing the (k+1)-th IMF:

$$IMF_{k+1}[n] = \frac{1}{I}\sum_{i=1}^{I} E_1(r_k[n] + \varepsilon_k E_k(w^i[n]))$$

Steps 5 and 6 are repeated until the residue becomes a monotonic function that cannot be further decomposed by EMD.

Step 7: Final Decomposition Result

The process concludes with the complete decomposition of the original signal:

$$x[n] = \sum_{k=1}^{K} IMF_k[n] + R[n]$$

where K represents the total number of IMFs and $R[n]$ is the final residue.

CEEMDAN offers several key advantages over its predecessors. By adding adaptive noise at each decomposition stage, the method significantly reduces mode mixing and maintains component purity. Furthermore, CEEMDAN demonstrates greater computational efficiency than EEMD, as it does not require a substantial increase in ensemble numbers to achieve accurate results.

In the context of exchange rate forecasting, the CEEMDAN method has proven particularly effective when integrated into hybrid forecasting frameworks. The typical workflow involves decomposing the exchange rate time series into a set of intrinsic mode functions (IMFs) using CEEMDAN to address noise and non-linear, non-stationary characteristics. Long Short-Term Memory (LSTM) networks are then applied to forecast each component individually, and the final prediction is obtained by aggregating the forecasts of all decomposed components.

3.3 Support Vector Machine Methodology

Support Vector Machines represent a supervised machine learning approach utilised for solving classification and regression tasks. Credit for developing this methodology belongs to scientist Vapnik in 1995 through his contributions to statistical learning theory. When this approach is applied to regression problems, it is termed Support Vector Regression, which seeks to estimate the values of the dependent variable y′ based on a regression function applied to a training dataset, where values belong to real numbers, such that L denotes the training sample size while D refers to the input space dimensions (García Nieto et al., 2018).

The regression function for the training dataset is mathematically formulated as follows:

$$f(x) = \langle w, x \rangle + b$$

This equation represents the dot product between the weight vector w and the input vector x, while the last variable refers to the intercept or bias coefficient in the model. The core concept of Support Vector Machine regression relies on employing a "loss function" that takes the value zero when the difference between the predicted value and the actual observed value falls within a range defined by the parameter ε, meaning this condition is satisfied. The region defined by this condition for all values is called the "ε-insensitive tube" (García Nieto et al., 2018). Data points that fall outside the boundaries of this tube are subject to penalties through slack variables that are determined based on their position relative to the tube. These constraints are mathematically formulated as follows:

$$\begin{gathered} y_i - \langle w, x_i \rangle - b \leq \varepsilon + \xi_i \\ \langle w, x_i \rangle + b - y_i \leq \varepsilon + \xi_i^* \\ \xi_i, \xi_i^* \geq 0 \end{gathered}$$

Accordingly, the loss function for Support Vector Regression is written in the following mathematical form:

$$\min_{w,b,\xi,\xi^*} \frac{1}{2} \| w \|^2 + C \sum_{i=1}^{L} (\xi_i + \xi_i^*)$$

where C represents the regularisation parameter or cost constant, while the other variables constitute the slack variables. The parameter C plays a crucial role in achieving balance between model flatness (margin) and the magnitude of training errors (slack variables).

To reach the optimal solution and minimise the error function, the Karush-Kuhn-Tucker optimality conditions are adopted. These are first-order mathematical conditions necessary for reaching the optimal solution in nonlinear programming problems with inequality constraints. By introducing Lagrange multipliers for all values, the problem is transformed into a Quadratic Programming problem that can be solved using the dual Lagrangian method (García Nieto et al., 2018). In the case of nonlinear data, input vectors from a low-dimensional space are mapped to a high-dimensional linear feature space through a nonlinear transformation. In this context, the nonlinear regression function is formulated as follows:

$$f(x) = \langle w, \phi(x) \rangle + b$$

Among the most commonly used kernel functions that satisfy Mercer's condition are:

Radial Basis Function (RBF) kernel:

$$K(x_i, x_j) = exp\left(-\frac{\| x_i - x_j \|^2}{2\sigma^2} \right)$$

Polynomial kernel:

$$K(x_i, x_j) = \left(a\langle x_i, x_j \rangle + b \right)^d$$

Sigmoid kernel:

$$K(x_i, x_j) = tanh\left(a\langle x_i, x_j \rangle + b \right)$$

In these equations, the variables represent parameters that determine the characteristics and behavior of the kernel function. The optimal combination of these hyperparameters is selected through applying the Grid Search technique, which is a comprehensive search strategy within a predefined subset space of the algorithm's hyperparameters, using Cross-Validation technique to select parameters that achieve the highest possible accuracy (García Nieto et al., 2018).

3.4 Long Short-Term Memory (LSTM)

Predicting financial time series one step ahead requires more than just the latest data; it also depends on incorporating historical information from previous time periods. Thanks to its self-feedback mechanism in the hidden layer, the Recurrent Neural Network (RNN) model offers significant advantages when dealing with long-term dependency issues. The LSTM unit comprises a memory cell that stores information, with updates managed through three specialised gates: the input gate, the forget gate, and the output gate. At any time step *t*, the fundamental variables include: x_t for the input data entering the cell, h_{t-1} for the previous time step's output, c_t for the memory cell's current value, and h_t for the cell's final output. Figure 1 illustrates the internal architecture of an LSTM unit, demonstrating how information flows through the memory cell and how the three control gates interact to regulate this flow.

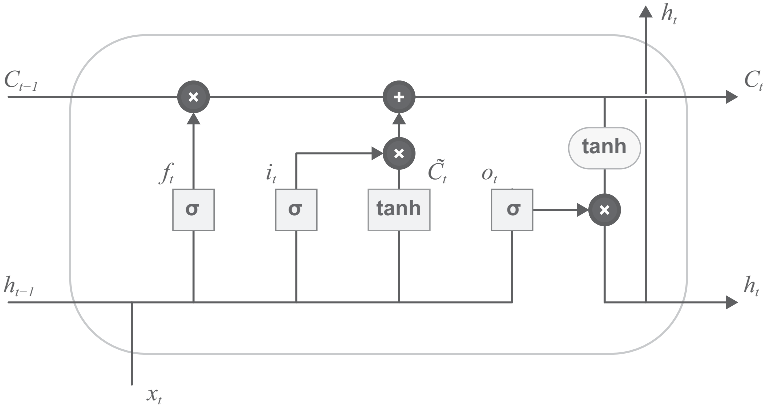

Figure 1: Architecture of LSTM Unit
Source: (Cao et al., 2019)

The LSTM unit's computational process unfolds through several sequential stages. First, the candidate memory cell value $\tilde{c}_t$ is calculated using the weight matrix W_c and bias vector b_c, expressed as: c_t=tanh(W_c·[h_{t-1},x_t]+b_c). Next, the input gate value i_t is determined through the sigmoid function, which regulates how current input data influences the memory cell state: i_t=σ(W_i·[h_{t-1},x_t]+b_i), where W_i represents the weight matrix and b_i the bias vector. The forget gate value f_t is then computed to control how historical information affects the memory cell state: f_t=σ(W_f·[h_{t-1},x_t]+b_f), with W_f as the weight matrix and b_f as the bias vector (Gers et al., 2000). As illustrated in Figure 1, the current memory cell value c_t is subsequently updated through the equation: c_t=f_t*c_{t-1}+i_t*$\tilde{c}_t$, where "*" denotes element-wise multiplication. This update relies on both the previous state c_{t-1} and the candidate cell value, regulated by the input and forget gates. The output gate value o_t is then calculated to determine what information from the memory cell state should be output: o_t=σ(W_o·[h_{t-1},x_t]+b_o), where W_o is the weight matrix and b_o is the bias vector. Finally, the LSTM unit produces its output through: h_t=o_t*tanh(c_t). These three control gates and the memory cell enable the LSTM model to efficiently store, retrieve, reset, and update long-term information. Moreover, the internal parameter sharing mechanism allows researchers to control output dimensions by adjusting the weight matrix dimensions. The architecture creates a substantial temporal delay between inputs and feedback, which prevents both gradient explosion and vanishing because the memory cell's internal state maintains a continuous error flow throughout the network (Cao et al., 2019).

3.5 Hybrid CEEMDAN-LSTM Model

Financial time series data inherently exhibit non-linear, non-stationary, and stochastic characteristics that pose significant forecasting challenges. Traditional single prediction models often struggle to capture the complex dynamics embedded within such data due to their volatility, irregular fluctuations, and multi-scale temporal patterns. Recognising these limitations, researchers have developed hybrid forecasting frameworks that integrate signal decomposition techniques with deep learning architectures. Specifically, combining Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN) and Long Short-Term Memory (LSTM) networks has emerged as a particularly effective approach for financial time series prediction. The fundamental premise underlying this hybrid methodology rests on transforming a complex forecasting problem into multiple simpler sub-problems through decomposition, where each component can be modeled more effectively than the original signal (Lin et al., 2021; Lin et al., 2022).

The first stage of the CEEMDAN-LSTM framework involves decomposing the original financial time series into several Intrinsic Mode Functions (IMFs) with different time scales plus one residual component. CEEMDAN represents an advanced evolution of the original Empirical Mode Decomposition (EMD) method, specifically designed to address the mode mixing problem that plagued earlier decomposition techniques. The fundamental principle of CEEMDAN involves adding adaptive white noise at each stage of the decomposition process, which serves to uniformly distribute different scales of the signal across appropriate IMF components. This adaptive noise addition distinguishes CEEMDAN from its predecessors, including the original EMD and even the Ensemble EMD (EEMD), by adding specific noise realisations at each decomposition stage rather than maintaining the same noise throughout the entire process, resulting in a more complete and stable decomposition with significantly reduced reconstruction error (Lin et al., 2021).

Mathematically, given an original financial time series $X(t)$, the CEEMDAN algorithm decom-poses it through a systematic process. Initially, white noise $\varepsilon^i(t)$ with a specific amplitude is added to the original signal, creating an ensemble of noisy signals: $X^i(t)=X(t)+\varepsilon_0 E_1(\varepsilon^i(t))$, where $i=1,2,\ldots,I$ represents the ensemble member index, $E_1(\cdot)$ denotes the operator that extracts the first EMD mode, and ε_0 represents the noise standard deviation coefficient. Each ensemble member undergoes EMD decomposition to obtain its first mode. The first IMF is calculated as the average across all ensemble members: $IMF_1(t) = \frac{1}{I}\sum_{i=1}^{I} IMF_1^i(t)$. Subsequently, the first residual is computed as: $r_1(t)=X(t)–IMF_1(t)$. For subsequent modes k=2,3,...,K, the algorithm adds the noise realisation $E_k(\varepsilon_i(t))$ to the residual $r_{k\text{-}1}(t)$, gen-erating: $r_{k\text{-}1}(t)+\varepsilon_{k\text{-}1}E_k\varepsilon_i(t)$. The k-th mode is obtained by averaging: $IMF_k(t) = \frac{1}{I}\sum_{i=1}^{I} E_1(r_{k-1}(t) + \varepsilon_{k-1}E_k(\varepsilon^i(t)))$, and the corresponding residual becomes: $r_k(t)=r_{k\text{-}1}(t)–IMF_k(t)$. This iterative procedure continues until the residual $r_K(t)$ can no longer be decomposed, at which point it becomes the final residual component $R(t)$. The original signal can then be perfectly reconstructed through summation: $X(t) = \sum_{k=1}^{K} IMF_k(t) + R(t)$ (Lin et al., 2022).

The decomposed IMFs represent distinct frequency characteristics of the original financial time series, with each IMF capturing a unique oscillatory mode. High-frequency IMFs typically capture short-term market volatility, rapid price fluctuations, and noise-like components that reflect immediate market reactions and trading dynamics. Medium-frequency IMFs reflect cyclical patterns and periodic behaviors associated with business cycles and seasonal effects. Low-frequency IMFs capture longer-term trends and structural changes in the market. The residual component represents the fundamental long-term trend or drift inherent in the data. This multi-resolution decomposition effectively transforms the non-stationary original series into several relatively stationary sub-series, each exhibiting more regular and predictable pattern. The key advantage of CEEMDAN over traditional EMD lies in its ability to eliminate mode mixing where oscillations of different scales inappropriately appear in the same IMF and to provide a complete decomposition with minimal reconstruction error, thereby ensuring that each characteristic time scale is represented by a unique and meaningful IMF component (Lin et al., 2021; Lin et al., 2022).

Following the decomposition stage, individual LSTM prediction models are established for each IMF component and the residual term. This approach ensures that the historical data's effects on prediction results are appropriately captured for each characteristic frequency component. The CEEMDAN-LSTM model utilises the decomposed data to obtain prediction sequences for each component separately, recognising that different frequency components exhibit distinct temporal patterns that require specialised modeling approaches. The theoretical justification for applying LSTM networks to each decomposed component stems from LSTM's inherent capability to learn long-term dependencies through its memory cell and gating mechanisms. However, in the context of the hybrid CEEMDAN-LSTM model, applying LSTM to decomposed components offers substantial advantages beyond those provided by applying LSTM directly to the original series (Lin et al., 2021).

Each decomposed IMF exhibits considerably more regular oscillatory behaviour and reduced non-stationarity compared to the original signal, making the learning task for individual LSTM networks substantially more tractable. Lin et al. (2022) emphasise that the modified ensemble empirical mode decomposition technique decomposes the original time series into several com-ponents with different frequencies, and each component is predicted separately by the LSTM neural network, which captures the long-term dependencies and complex temporal patterns more effectively. The LSTM network for each component is trained independently using a supervised learning framework. For the *k*-th IMF component $IMF_k(t)$, the input sequence consists of historical values $[IMF_k(t-n), IMF_k(t-n+1),...,IMF_k(t-1)]$, where n represents the lookback window size, and the target output is $IMF_k(t)$. The LSTM network learns to map the historical patterns in each frequency component to future values by adjusting its internal parameters, spe-cifically, the weight matrices and bias vectors for the input gate, forget gate, and output gate, as well as the candidate memory cell, through the backpropagation through time (BPTT) algorithm, minimising a loss function such as mean squared error (MSE): $L = \frac{1}{N}\sum_{t=1}^{N}(IMF_k(t) - \widehat{IMF_k}(t))^2$, where $\widehat{IMF_k}(t)$ represents the LSTM's prediction and N denotes the number of training samples (Lin et al., 2021; Lin et al., 2022).

The architecture of each LSTM network can be customised according to the characteristics of its corresponding IMF component. Typically, high-frequency IMFs require LSTM networks with fewer layers and hidden units due to their simpler and more regular patterns, while low-frequency IMFs and the residual component may benefit from deeper networks with more hidden units to capture their complex long-term dependencies and intricate temporal structures. The training process for each LSTM model is conducted independently, allowing for component-specific hyperparameter optimisation. This modular approach provides considerable flexibility in model design and enables parallel computation during both training and prediction phases, potentially reducing overall computational time compared to monolithic architectures. Furthermore, this decomposition-based strategy allows the model to allocate computational resources more efficiently, dedicating more complex network architectures to components that genuinely require them while using simpler structures for more straightforward patterns (Lin et al., 2022).

The final stage of the CEEMDAN-LSTM model involves reconstructing the predictions from all individual LSTM models to obtain the final forecast of the original time series. Given the predicted values $IMF_k(t)$ from the LSTM model for each k-th IMF component and $R(t)$ from the LSTM model for the residual component, the final prediction $X(t)$ is computed through straightforward summation: $X(t) = \sum_{k=1}^{K} IMF_k(t) + R(t)$. This linear re-construction formula directly follows from the decomposition property of CEEMDAN, where the original signal equals the sum of all its components. The theoretical elegance of this ensem-ble approach lies in its simplicity and interpretability: by accurately predicting each frequency component independently and then reconstructing them, the model effectively captures the mul-ti-scale temporal dynamics of the original financial time series without requiring complex en-semble weighting schemes, meta-learning procedures, or additional optimisation steps (Lin et al., 2021; Lin et al., 2022).

3.6 Evaluation Metrics

To evaluate the predictive capability of the proposed models, the study employed three standard performance indicators. The first of these is the Mean Absolute Error (MAE), which is defined by the equation below (Manowska, 2020):

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

The second metric is the Root Mean Square Error (RMSE), expressed by the equation:

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

The third metric is the Mean Absolute Percentage Error (MAPE), which determines the prediction accuracy as a percentage according to the following formula:

$$MAPE = \frac{1}{n}\sum_{i=1}^{n} \left| \frac{y_i - \hat{y}_i}{y_i} \right| 100\%$$

Here, y_i denotes the observed value and y_i represents the predicted value for the i-th sample, while n refers to the total number of data points in the testing set.

3.7 Statistical Significance, Robustness, and Bootstrap Prediction Intervals

Because error measures alone do not establish whether forecast differences are statistically meaningful, the analysis applies the Diebold–Mariano test of equal predictive accuracy under squared-error loss (Diebold & Mariano, 1995). Given the limited size of the primary test sample, the small-sample modified statistic of Harvey, Leybourne, and Newbold (1997) is reported with two-sided p-values.

Robustness is evaluated by repeating the chronological holdout exercise using an 80/20 split. This alternative test period covers 62 months and permits assessment of whether relative model ranking is stable over a longer out-of-sample horizon.

Finally, the 95% uncertainty bounds for the 12-month forecast are generated through a residual-bootstrap procedure rather than a normal-error assumption. Specifically, centered residuals from the primary CEEMDAN-LSTM holdout are resampled with replacement 5,000 times, and the 2.5th and 97.5th percentiles of the simulated forecasts define the reported prediction limits (Efron & Tibshirani, 1993).

4. RESULTS AND DISCUSSIONS

4.1 Preliminary Data Processing and Descriptive Analysis

This study relied on monthly data for the US dollar exchange rate against the Algerian dinar, spanning a period from January 2000 through October 2025, representing 310 consecutive monthly observations. Table (1) reveals the overall statistical characteristics of this data, where the average exchange rate throughout this period settled at 95.05 dinars per US dollar. However, the substantial standard deviation of 25.92 dinars indicates the presence of sharp fluctuations in currency value. Notably, the lowest recorded price was 61 dinars in the early stages of the study period, while the rate jumped to reach its peak at 145.57 dinars, representing a wide price gap of 84.57 dinars. This considerable range in price volatility clearly underscores the necessity of employing modern and sophisticated forecasting models capable of absorbing and handling this notable degree of market instability.

Table 1. Statistical Characteristics of the USD/DZD Exchange Rate (2000-2025)

Indicator	Mean	Std. Error	Median	Variance	Kurtosis	Skewness	Min	Max	Count
Value	95.05	1.47	79.16	671.87	-1.35	0.55	61.00	145.57	310

Source: Authors' elaboration using python 3.13

Through analysing the skewness and kurtosis coefficients presented in Table (1), we observe that the positive skewness coefficient (+0.55) indicates that the data distribution is skewed to the right, meaning that the majority of observations recorded prices lower than the arithmetic mean. However, the presence of periods that witnessed sharp increases in the exchange rate led to a noticeable elevation of the overall average. As for the negative kurtosis coefficient (-1.35), it suggests that the distribution is flat (Platykurtic), which means the data is widely spread across the price range rather than being concentrated around a specific central value. These statistical characteristics reflect the nature of the Algerian economic environment, which has been influenced by multiple internal and external factors, including oil price fluctuations, monetary policies, and global economic shocks.

Moreover, the substantial gap between the arithmetic mean (95.05) and the median (79.16) clearly demonstrates the impact of extreme values on the distribution. The lower median reflects the actual reality for most of the time period, while the mean is significantly affected by periods that witnessed a sharp decline in the Algerian dinar's value during recent years. This evident discrepancy emphasises the necessity of employing advanced analytical methods that don't rely solely on traditional central statistics, but rather take into account the dynamic structure and non-linear characteristics of the time series. This justifies the selection of the hybrid CEEMDAN-LSTM model in this study.

4.2 Hybrid CEEMDAN-LSTM Model Methodology

The data was divided into two sets: a training set comprising 279 observations (approximately 90% of the total data), and a testing set containing 31 observations (approximately 10%), while maintaining the natural chronological order of the data without random shuffling, to ensure valid evaluation within the time series context. This division is appropriate for long time series, as the testing set of 31 data points (equivalent to roughly 2.5 years) provides sufficient size to assess the model's ability to generalise and forecast accurately.

The analysis process begins by applying CEEMDAN to the original time series. The decomposition process produced five intrinsic mode functions (IMFs) in addition to a residual component, where each component represents a different frequency level in the original data, as illustrated in Figure (2). This multi-level decomposition helps separate short-term fluctuations and irregular shocks from medium- and long-term movements, making it easier for the LSTM models to learn from each pattern individually rather than handling the raw series as a single process.

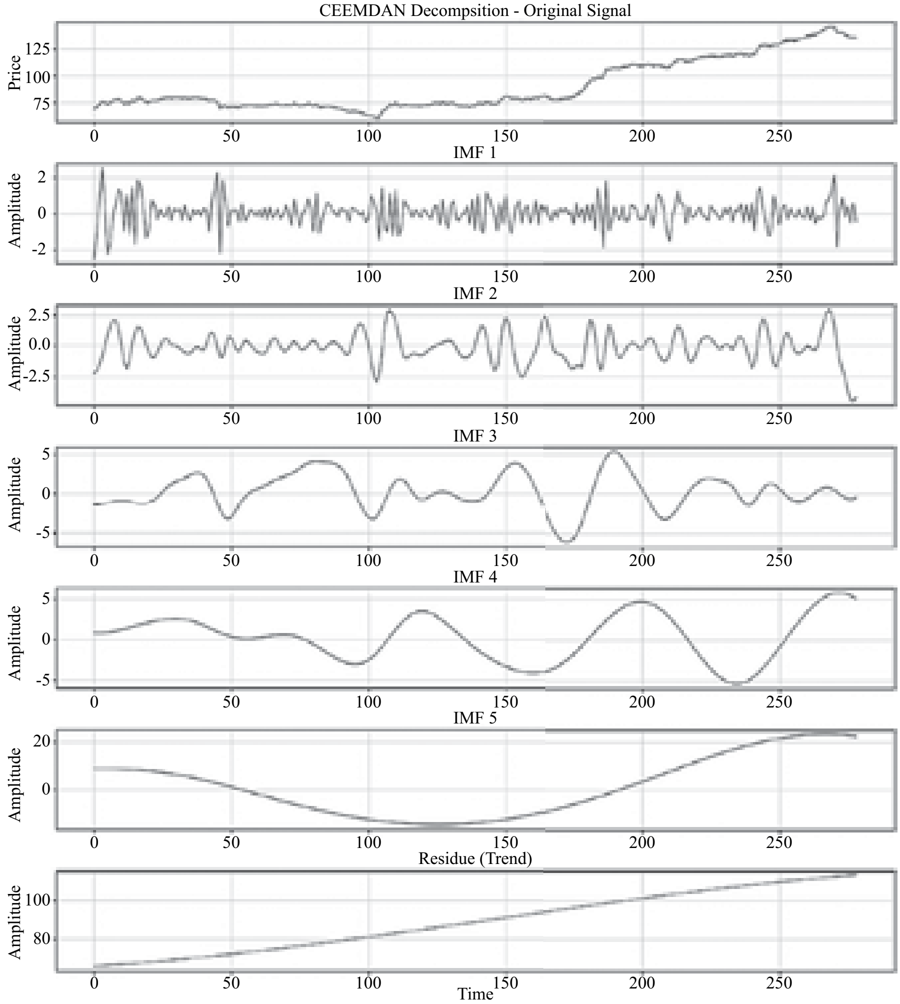

Figure 2: CEEMDAN decomposition of USD/DZD exchange rate into six frequency components
Source: Authors' elaboration using python 3.13

Figure (2) illustrates how the original time series was decomposed into multiple components. IMF1, IMF2, and IMF3 capture higher-frequency movements associated with short-term fluctuations, whereas IMF4, IMF5, and the residual component represent medium- and long-term dynamics. This separation allows the forecasting model to learn patterns at different horizons more effectively.

Each CEEMDAN component was normalised using a scaler fitted on the corresponding training segment only. The LSTM specification uses two stacked layers with 32 and 16 units, trained for 500 epochs with the Adam optimiser (learning rate = 0.001) and mean squared error loss. A parsimonious architecture was retained to limit over-parameterisation in the short monthly series; the same specification was maintained across the primary and robustness evaluations. The lookback window is two months for IMF1–IMF3 and four months for IMF4, IMF5, and the residual; the standalone LSTM uses four lags. The SVR benchmark uses an RBF kernel, with C, gamma, and epsilon selected by grid search under four-fold time-series cross-validation on the training segment only. A fixed random seed of 123 ensures reproducibility. The complete hyperparameter settings and selection procedures for all models are summarised in Table 2.

Table 2. Hyperparameter Specification and Selection Protocol

Model	Hyperparameter	Selected value	Selection protocol
CEEMDAN	Trials; noise amplitude	50; 0.005	Fixed reproducible specification
CEEMDAN-LSTM / LSTM	LSTM layers and units	2 layers: 32, 16	Parsimonious fixed specification
CEEMDAN-LSTM / LSTM	Optimiser; learning rate; loss	Adam; 0.001; MSE	Fixed specification
CEEMDAN-LSTM / LSTM	Epochs; dropout	500; 0.00	Held constant across splits
CEEMDAN-LSTM	Lookback window	2 for IMF1–IMF3; 4 otherwise	Frequency-based rule
Traditional LSTM	Lookback window	4	Held constant across splits
SVR	Kernel; C; gamma; epsilon	RBF; 100; 0.01; 0.001	Grid search + TimeSeriesSplit(4)

Source: Authors' specification and computations using Python 3.13.

4.3 Model Performance During Training

Table (3) reports the final training losses obtained from the six component-specific LSTM models and fully documented specification. The MSE loss function and Adam optimisation algorithm were applied uniformly to all components for 500 epochs.

The residual and IMF5 components exhibit the smallest final losses (0.0001 and 0.0004), reflecting the greater regularity of the lower-frequency dynamics. By contrast, IMF1 and IMF2 are more difficult to fit because they capture higher-frequency fluctuations. The training-loss curves in Figure (3) show declining loss profiles under the documented specification.

Table 3. Training Results of the Component-Specific LSTM Models

Component	Window Size	Final Loss	Number of Epochs	Characteristics
IMF 1	2	0.0188	500	High frequency (short-term fluctuations)
IMF 2	2	0.0132	500	High frequency
IMF 3	2	0.0032	500	Medium frequency
IMF 4	4	0.0027	500	Medium frequency
IMF 5	4	0.0004	500	Low frequency (trend)
Residue	4	0.0001	500	Long-run component

Source: Authors' computations using Python 3.13.

Figure (3) presents the training-loss curves for the six component-specific LSTM models.

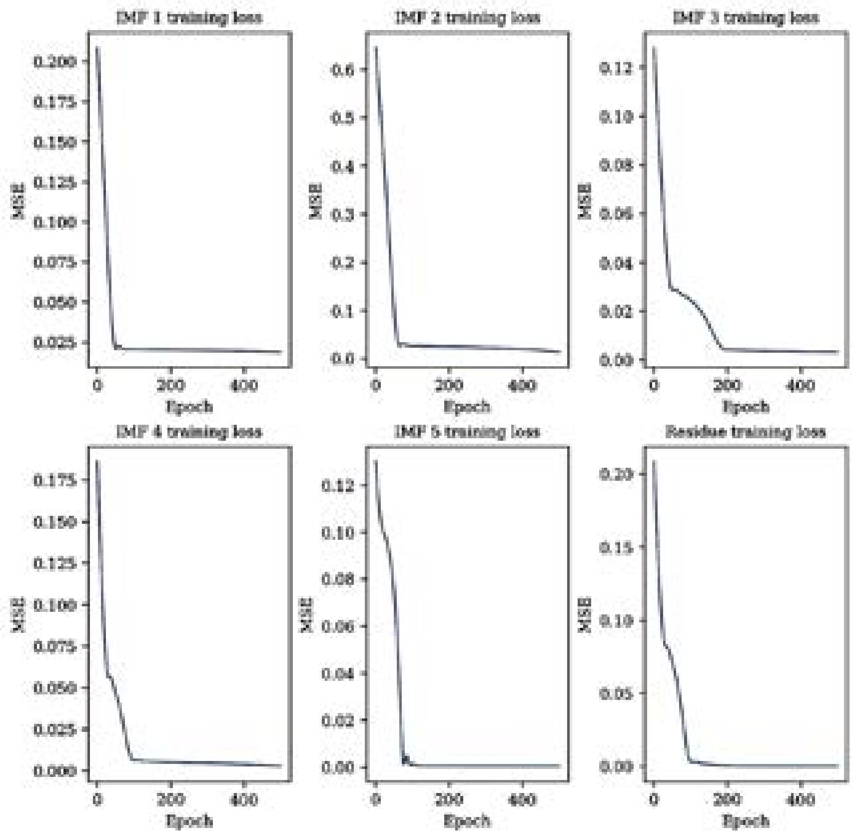

Figure 3: Training-Loss Curves for the Component-Specific LSTM Models
Source: Authors' computations using Python 3.13.

4.4 Model Performance Evaluation on Test Data

After estimation, the component forecasts were aggregated to reconstruct the CEEMDAN-LSTM forecast. Table (4) compares the out-of-sample results on the primary 90/10 chronological holdout with the standalone LSTM and the tuned SVR benchmark.

Table 4. Performance Comparison on the Primary 90/10 Holdout (31 Observations)

Model	MAE	RMSE	MAPE (%)
Traditional LSTM	1.4710	1.8070	1.1094
SVR (RBF)	0.8406	1.0769	0.6304
CEEMDAN-LSTM	0.6393	0.8149	0.4782

Source: Authors' computations using Python 3.13.

On the primary test period, CEEMDAN-LSTM yields the lowest errors on all three metrics. Its MAE of 0.6393 dinars is 56.5% lower than the standalone LSTM error and 23.9% lower than the tuned SVR error. The corresponding MAPE of 0.4782% indicates a close match to the observed USD/DZD values in this holdout.

As shown in Table 5, The modified Diebold–Mariano results support the primary-sample ranking: equal predictive accuracy is rejected at the 5% level against both benchmarks. This provides statistical evidence for the hybrid model's improvement over the primary evaluation window.

Table 5. Modified Diebold–Mariano Tests on the Primary 90/10 Holdout

Model comparison (squared-error loss)	MDM statistic	p-value	Decision at 5%
CEEMDAN-LSTM vs Traditional LSTM	-3.3369	0.0023	Reject equal accuracy; hybrid preferred
CEEMDAN-LSTM vs SVR (RBF)	-2.2930	0.0290	Reject equal accuracy; hybrid preferred

Source: Authors' computations; negative statistic indicates lower loss for CEEMDAN-LSTM.

The improvement is consistent with the decomposition rationale: CEEMDAN isolates components operating at different frequencies, while component-specific LSTM networks learn comparatively homogeneous patterns before their forecasts are recombined.

Figure (4) plots actual observations against the CEEMDAN-LSTM forecasts in the primary holdout period.

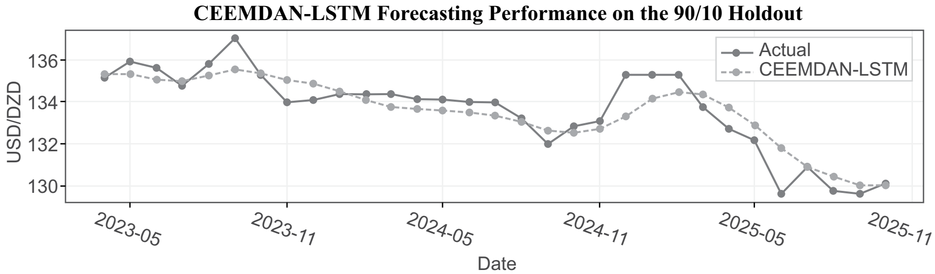

Figure 4: Actual and CEEMDAN-LSTM Forecast Values on the Primary 90/10 Holdout
Source: Authors' computations using Python 3.13.

Figure (4) shows that the hybrid forecast closely follows the observed movements in the primary holdout, although deviations remain at some local turning points.

4.5 Future Forecasting for Twelve Months

The CEEMDAN-LSTM model was re-estimated using the full series and used to generate recursive monthly forecasts from November 2025 through October 2026. In response to the non-normal distributional evidence in Table (1), Table (6) reports residual-bootstrap 95% prediction intervals based on 5,000 resamples rather than intervals obtained under a normal-error assumption.

Table 6. Future USD/DZD Forecasts with 95% Residual-Bootstrap Prediction Intervals

Month	Point Forecast (Dinar)	Lower 95% Bootstrap Bound	Upper 95% Bootstrap Bound
November 2025	130.37	128.12	132.30
December 2025	130.07	127.82	132.00
January 2026	129.85	127.59	131.78
February 2026	129.86	127.60	131.78
March 2026	129.95	127.70	131.88
April 2026	129.90	127.65	131.83
May 2026	129.92	127.67	131.85
June 2026	129.98	127.73	131.91
July 2026	130.05	127.80	131.98
August 2026	130.08	127.82	132.00
September 2026	130.14	127.89	132.07
October 2026	130.21	127.96	132.14

Source: Authors' computations using 5,000 centered residual-bootstrap replications in Python 3.13.

The forecast path indicates relative stability around 130 dinars per US dollar. The point forecast declines from 130.37 in November 2025 to 129.85 in January

2026 and then rises gradually to 130.21 in October 2026, implying only a modest change over the forecast horizon.

The bootstrap limits are asymmetric around the point estimates and avoid imposing normality on future errors. For example, the 95% interval for November 2025 is [128.12, 132.30], while the October 2026 interval is [127.96, 132.14]. These uncertainty bounds should be interpreted as baseline ranges rather than protection against structural breaks or exceptional policy and commodity-price shocks.

Figure (6) presents recent observations, primary-holdout predictions, and the twelve-month forecast path with the residual-bootstrap prediction interval.

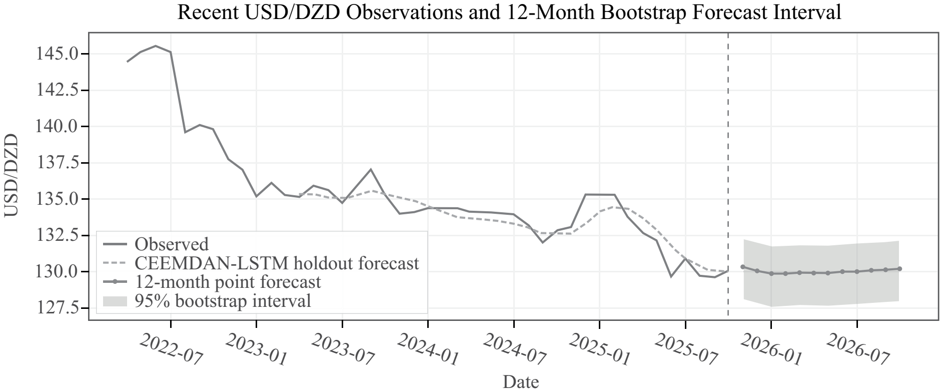

Figure 6: Recent USD/DZD Observations and Twelve-Month Forecast with Bootstrap Interval

Source: Authors' computations using Python 3.13.

The forecast band in Figure (6) makes explicit the uncertainty surrounding the otherwise stable point forecast and provides a more defensible basis for risk analysis than normality-based bounds.

4.6 Discussions and Policy Implications

The results provide evidence in favour of decomposition-based forecasting for the USD/DZD exchange rate. In the primary evaluation, the CEEMDAN-LSTM model achieves lower forecasting errors than both the standalone LSTM and the tuned SVR benchmarks. Furthermore, the modified Diebold–Mariano tests indicate that these improvements are statistically significant, providing empirical support for the forecasting advantage of the proposed hybrid framework.

From an economic perspective, the decomposition results remain meaningful. High-frequency intrinsic mode functions (IMFs) may reflect short-term responses to temporary shocks, market sentiment, and transitory trade or policy disturbances. In contrast, lower-frequency IMFs and the residual component capture more persistent dynamics that may be associated with hydrocarbon export revenues, import demand, inflation differentials, monetary conditions, and exchange-rate management policies. By separating these heterogeneous movements into more homogeneous components, CEEMDAN facilitates more effective pattern learning by the LSTM networks and contributes to improved forecasting performance.

These findings have several policy implications. First, more accurate short-term forecasts can support the timing of foreign-exchange interventions, liquidity planning, and reserve management by monetary authorities. Second, reliable medium-term exchange-rate projections can assist public institutions and import-dependent firms in budgeting, cash-flow management, and hedging decisions. Third, although the twelve-month forecast suggests relative stability in the USD/DZD exchange rate, it should be interpreted as a baseline scenario rather than a guarantee of future stability. Policymakers and market participants should complement baseline forecasts with scenario analysis to assess the potential effects of adverse developments such as oil-price fluctuations, inflationary pressures, or external financial shocks.

Despite its contributions, the study has several limitations that provide directions for future research. The analysis remains univariate and focuses on a single currency pair, while structural changes and exogenous shocks are incorporated only indirectly through the historical exchange-rate series. Future research could extend the framework by incorporating additional macroeconomic variables, including oil prices, inflation rates, interest-rate differentials, and foreign-exchange reserves. Moreover, further comparisons with advanced deep-learning approaches, such as attention-based models and transformer architectures, may provide additional insights into the relative strengths of decomposition-based forecasting methods in exchange-rate prediction.

5. CONCLUSIONS

Overall, the findings demonstrate the effectiveness of the CEEMDAN-LSTM hybrid framework for forecasting the USD/DZD exchange rate using 310 monthly observations covering the period from January 2000 to October 2025. In the primary evaluation, the hybrid model achieved MAE = 0.6393, RMSE =

0.8149, and MAPE = 0.4782%, outperforming both the standalone LSTM and the tuned SVR benchmarks. Furthermore, the modified Diebold–Mariano tests confirmed that these forecasting improvements were statistically significant relative to both competing models.

The prediction intervals generated through the residual bootstrap approach suggest a relatively stable baseline trajectory for the USD/DZD exchange rate over the subsequent twelve months, fluctuating around 130 Algerian dinars per US dollar while explicitly accounting for forecast uncertainty. These results indicate that the CEEMDAN-LSTM framework can serve as a useful tool for exchange-rate forecasting, financial planning, and exchange-rate risk management.

Despite these encouraging findings, practical applications of the model should be accompanied by regular model updating and continuous monitoring of new market information. In addition, forecast-based decisions should be complemented by scenario analysis to account for potential economic and financial shocks. Future research may further enhance forecasting performance by incorporating relevant macroeconomic variables and comparing the proposed framework with more advanced deep-learning architectures.

ACKNOWLEDGEMENTS

The authors thank the editor and the anonymous reviewers for their constructive comments, which helped improve the manuscript.

Conflict of interests

The authors declare there is no conflict of interest.

REFERENCES

Adesina, O. S., & Obokoh, L. O. (2025). A hybrid framework of deep learning and traditional time series models for exchange rate prediction. *Scientific African, 29*, Article e02818. https://doi.org/10.1016/j.sciaf.2025.e02818

Al-Gounmeein, R. S., & Ismail, M. T. (2020). Forecasting the exchange rate of the Jordanian dinar versus the US dollar using a Box-Jenkins seasonal ARIMA model. *International Journal of Mathematics and* Computer *Science, 15*(1), 27–40. http://ijmcs.future-in-tech.net

Amri, I. F., Yunanita, N., Lestari, F. A., & Dhani, O. R. (2025). Rupiah exchange rate prediction against the US dollar using a deep neural network with a multi-output sliding window approach. MethodsX. Advance online publication. https://doi.org/10.1016/j.mex.2025.103692

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

Darvas, Z., & Schepp, Z. (2025). Forecasting the daily exchange rate of the UK pound sterling against the US dollar. *Finance Research Letters, 71*, Article 106451. https://doi.org/10.1016/j.frl.2024.106451

De Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, *22*(3), 443–473. https://doi.org/10.1016/j.ijforecast.2006.01.001

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, *13*(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

Efron, B., & Tibshirani, R. J. (1993). An introduction to the bootstrap. Chapman & Hall/CRC. https://doi.org/10.1007/978-1-4899-4541-9

Fan, W., Song, M., & Du, Z. (2021). Forecasting carbon price based on CEEMDAN and LSTM optimized by whale optimization algorithm. *Journal of Forecasting*, *40*(7), 1235–1250. https://doi.org/10.1002/for.2761

Flores-Sosa, M., León-Castro, E., Merigó, J. M., & Yager, R. R. (2022). Forecasting the exchange rate with multiple linear regression and heavy ordered weighted average operators. *Knowledge-Based Systems, 248*, Article 108863. https://doi.org/10.1016/j.knosys.2022.108863

García, F., Guijarro, F., Oliver, J., & Tamošiūnienė, R. (2024). Foreign exchange forecasting models: LSTM and BiLSTM comparison. *Engineering Proceedings, 68*(1), Article 19. https://doi.org/10.3390/engproc2024068019

García Nieto, P.J., Sánchez Lasheras, F., García-Gonzalo, E., & de Cos Juez, F.J. (2018). PM_{10} concentration forecasting in the metropolitan area of Oviedo (Northern Spain) using models based on SVM, MLP, VARMA and ARIMA: A case study. *Science of the Total Environment*, 621, 753-761. https://doi.org/10.1016/j.scitotenv.2017.11.291

Gers, F. A., Schmidhuber, J., & Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. *Neural Computation*, *12*(10), 2451-2471. https://doi.org/10.1162/089976600300015015

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, *13*(2), 281–291. https://doi.org/10.1016/S0169-2070(96)00719-4

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, *9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Huang, N. E., Shen, Z., & Long, S. R. (1999). A new view of nonlinear water waves: The Hilbert spectrum. *Annual Review of Fluid Mechanics, 31*(1), 417-457. https://doi.org/10.1146/annurev.fluid.31.1.417

Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., & Liu, H. H. (1998). The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society A:*

Mathematical, Physical and Engineering Sciences, *454*(1971), 903–995. https://doi.org/10.1098/rspa.1998.0193

Kamouh, M. M. N. (2026). Forecasting exchange rates using artificial neural networks: An applied study on the Arab Republic of Egypt. *Arab Journal of Administration, 46*(6), 1–22. https://doi.org/10.21608/aja.2024.331953.1740

Li, X. (2006). Temporal structure of neuronal population oscillations with empirical mode decomposition. *Physics* Letters *A*, *356*(3), 237-241. https://doi.org/10.1016/j.physleta.2006.03.045

Lin, C.-S., Chiu, S.-H., & Lin, T.-Y. (2012). Empirical mode decomposition–based least squares support vector regression for foreign exchange rate forecasting. *Economic Modelling, 29*(6), 2583–2590. https://doi.org/10.1016/j.econmod.2012.07.018

Lin, Y., Yan, Y., Xu, J., Liao, Y., & Ma, F. (2021). Forecasting stock index price using the CEEMDAN-LSTM model. *The North American Journal of Economics and Finance*, 57, 101421. https://doi.org/10.1016/j.najef.2021.101421

Lin, Y., Liao, Q., Lin, Z., Tan, B., & Yu, Y. (2022). A novel hybrid model integrating modified ensemble empirical mode decomposition and LSTM neural network for multi-step precious metal prices prediction. *Resources Policy*, 78, 102884. https://doi.org/10.1016/j.resourpol.2022.102884

Liu, P., Wang, Z., Liu, D., Wang, J., & Wang, T. (2023). A CNN-STLSTM-AM model for forecasting USD/RMB exchange rate. *Journal of Engineering Research, 11*(3), Article 100079. https://doi.org/10.1016/j.jer.2023.100079

Lyócsa, Š., Plíhal, T., & Výrost, T. (2024). Forecasting day-ahead expected shortfall on the EUR/USD exchange rate: The (I)relevance of implied volatility. *International Journal of Forecasting, 40*(4), 1275–1301. https://doi.org/10.1016/j.ijforecast.2023.11.003

Manowska, A. (2020). Using the LSTM Network to Forecast the Demand for Electricity in Poland. Applied Sciences, 10(23), 8455. https://doi.org/10.3390/app10238455

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

Ren, Y., Suganthan, P. N., & Srikanth, N. (2016). Ensemble methods for wind and solar power prediction - A state-of-the-art review. *Renewable and Sustainable Energy Reviews*, *50*, 82–91. https://doi.org/10.1016/j.rser.2015.04.081

Torres, M. E., Colominas, M. A., Schlotthauer, G., & Flandrin, P. (2011). A complete ensemble empirical mode decomposition with adaptive noise. *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4144–4147. https://doi.org/10.1109/ICASSP.2011.5947265

Vapnik, V.N. (1995). The Nature of Statistical Learning Theory. Springer, New York. https://doi.org/10.1007/978-1-4757-2440-0

Wu, Z., & Huang, N. E. (2009). Ensemble empirical mode decomposition: A noise-assisted data analysis method. *Advances in Adaptive Data Analysis*, *1*(1), 1–41. https://doi.org/10.1142/S1793536909000047

Yu, L., Wang, S., & Lai, K. K. (2008). Forecasting crude oil price with an EMD-based neural network ensemble learning paradigm. *Energy Economics*, *30*(5), 2623-2635. https://doi.org/10.1016/j.eneco.2008.05.003

Zhang, B. (2018). Foreign exchange rates forecasting with an EMD-LSTM neural networks model. Journal *of Physics: Conf. Ser. 1053 012005*. https://doi.org/10.1088/1742-6596/1053/1/012005

ЦЕЕМДАН-ЛСТМ МОДЕЛ ЗА ПРЕДВИЂАЊЕ ОДНОСА КУРСА АМЕРИЧКОГ ДОЛАРА И АЛЖИРСКОГ ДИНАРА

1 Абделкадар Сехд, Универзитетски центар у Магнији, Економски факултет, Алжир
2 Мохамед Мекидиш, Универзитетски центар у Магнији, Економски факултет, Алжир
3 Хасин Кахуи, Универзитетски центар у Магнији, Економски факултет, Алжир

САЖЕТАК

Овај рад развија хибридни модел за предвиђање девизног курса америчког долара према алжирском динару комбиновањем ЦЕЕМДАН декомпозиције са ЛСТМ мрежама како би се ријешили изазови високе волатилности и сложених временских зависности на валутним тржиштима. Модел је анализирао 310 историјских опажања са високом волатилношћу (стандардна девијација: 25,92). ЦЕЕМДАН је декомпоновао временску серију на пет фреквентних компоненти и резидуални дио, при чему је свака компонента обрађена независном ЛСТМ мрежом специјализованом за учење образаца унутар свог временског опсега. ЦЕЕМДАН-ЛСТМ модел је постигао најниже вриједности грешке, са МАПЕ од 0,56%, и надмашио традиционалне ЛСТМ и СВМ моделе. Прогноза за 12 мјесеци указује на релативну стабилност девизног курса са благим падом од око 0,48%. Резултати показују да декомпозиција оригиналне серије прије СЛТМ моделирања побољшава тачност предвиђања раздвајањем краткорочног шума од средњорочних и дугорочних кретања. Ова методологија може послужити финансијским институцијама и доносиоцима економских одлука као користан алат за управљање ризиком, стратегије заштите и краткорочно и средњорочно планирање.

Кључне ријечи: *ЦЕЕМДАН, ЛСТМ, предвиђање девизног курса, финансијске временске серије, дубоко учење, управљање ризиком.*