Journal of Innovative Applied Mathematics and Computational Sciences

J. Innov. Appl. Math. Comput. Sci. **5**(2) (2025), 291–311. DOI: 10.58205/jiamcs.v5i2.1929

http://jiamcs.univ-mila.dz/

Forecasting Algerian time series: a comparative study of ANN and SARIMA models

Khaoula Djaber and Mouna Merzougui

LaPS Laboratory, Badji Mokhtar Annaba University, P.O. Box 12, 23000 Annaba, Algeria.

Received May 7, 2025, Accepted December 23, 2025, Published January 08, 2026

Abstract. Accurate time series forecasting is essential for informed decision-making in economic planning, financial management, and environmental monitoring. Traditional linear models such as the Seasonal Autoregressive Integrated Moving Average (SARIMA) are widely used but often fail to capture the nonlinear and complex dynamics inherent in many real-world datasets. In recent years, Artificial Neural Networks (ANNs) have emerged as a powerful alternative, capable of modeling such complexities without relying on rigid assumptions. This study applies ANN models to three Algerian time series: Gross Domestic Product (GDP), the US Dollar Algerian Dinar (USD/DZD) exchange rate, and monthly average temperature. The forecasting performance of ANN models is benchmarked against SARIMA models. Empirical results demonstrate that ANNs consistently outperform SARIMA models in terms of predictive accuracy across all datasets, highlighting their robustness and adaptability in diverse forecasting contexts.

Keywords: Artificial Neural Networks, SARIMA models, GDP, temperature, exchange rate.

2020 Mathematics Subject Classification: 62M10, 68T05, 91B84.

1 Introduction

Time series forecasting plays a pivotal role across various disciplines, serving as a foundational tool for strategic planning in economics, finance, and climate science [7]. In Algeria, reliable forecasts of key indicators such as GDP, exchange rates, and temperature trends are vital for economic policy development, currency risk management, and climate-related planning. Traditional statistical approaches, most notably the Seasonal Autoregressive Integrated Moving Average (SARIMA) model, have long been standard tools for modeling and forecasting time series data. While these models are well-suited for capturing linear and seasonal structures, their ability to handle nonlinear dynamics remains limited.

To overcome these limitations, recent research has increasingly turned to machine learning methods, particularly Artificial Neural Networks (ANNs), which offer greater flexibility in modeling complex, nonlinear patterns [2,5]. ANNs do not require strict assumptions about

✉Corresponding author. Email: mouna.merzougui@univ-annaba.dz

ISSN (electronic): 2773-4196

the underlying data-generating process and are capable of approximating highly nonlinear relationships, making them particularly attractive for real-world forecasting applications. Foundational concepts and architectures of deep learning are comprehensively presented in the seminal works of [2], [9], and [5], providing theoretical grounding for the development and application of ANN models in real-world problems.

Numerous studies have demonstrated the effectiveness of ANN-based models in diverse forecasting contexts. [10] proposed a multi-model ANN approach for mid- and long-term load forecasting, achieving significant improvements over traditional methods. [4] showed that ANN models outperformed other classical time series and machine learning methods for forecasting GDP growth rates. Similarly, [11] reported that a nonlinear autoregressive neural network (NAR) provided more accurate rainfall forecasts than traditional ARMA and SETAR models. [8] found that machine learning techniques, including recurrent neural networks (RNN) and long short-term memory (LSTM) models, achieved better forecasting performance than ARIMA for exchange rate prediction. [6] also showed that ANN models offered superior forecasting accuracy compared to ARIMA models when applied to macroeconomic indicators such as inflation, exchange rates, and GDP. [13] demonstrated that ANN-based models produced significantly more realistic inflation forecasts than those found in the official budget law and Medium-Term Program of Türkiye, thereby highlighting the model's practical relevance for national economic planning.

Although such comparative analyses between ANN and traditional models (such as ARIMA or SARIMA) have been widely conducted in other countries, very few studies have explored this question within the Algerian context. Most existing Algerian contributions have focused on narrow domains, such as the use of ANNs for exchange rate forecasting [1]. However, a comprehensive, cross-sectoral comparison of ANN and SARIMA models applied to multiple Algerian time series (economic, financial, and climatic) remains largely absent from the literature. This study aims to fill this gap by providing empirical evidence on the relative forecasting performance of these two approaches using real Algerian data.

Specifically, we investigate the effectiveness of Neural Network Autoregressive (NNAR) models feedforward neural networks adapted for time series forecasting in predicting three representative Algerian time series from diverse domains: economic activity (GDP), financial markets (USD/DZD exchange rate), and environmental data (monthly average temperature). The forecasting performance of the ANN models is compared to that of SARIMA models using standard accuracy metrics. The results provide compelling evidence of the superior ability of ANN-based approaches to capture the underlying complexity of these time series, supporting their practical adoption for forecasting applications within the Algerian context. The findings of this study hold significant policy relevance: improved GDP forecasts can enhance economic planning, accurate exchange rate predictions can assist in managing currency risks, and reliable temperature forecasts can support environmental and climate adaptation strategies.

The remainder of this paper is organized as follows. Section 2 outlines the methodological framework, including data preprocessing and a presentation of the SARIMA and ANN models. Section 3 illustrates the empirical application to Algerian time series data. Finally, Section 4 concludes the paper by summarizing the key findings and highlighting potential directions for future research.

2 Methodology

2.1 Data and preprocessing

This study focuses on three Algerian time series datasets, each representing a distinct domain:

- **Gross Domestic Product:** Annual GDP data for Algeria from 1960 to 2023 (in billions of USD), obtained from Macrotrends (https://www.macrotrends.net/global-metrics/countries/DZA/algeria/population). This series reflects the country's economic trajectory, influenced by oil price fluctuations, structural reforms, and demographic changes. A strong upward trend is observed. The test set comprises the most recent years from 2016 to 2023.

- **Exchange Rate (USD/DZD):** Monthly exchange rate data for the Algerian Dinar (DZD) against the US Dollar (USD), spanning from January 1, 2022, to January 1, 2024. Data were retrieved from Yahoo Finance and reflect short-term financial market dynamics. The final four months (from September 28, 2023, onward) were reserved for testing.

- **Temperature:** Monthly average temperature data for Algeria, from January 1991 to December 2016, sourced from the World Bank database. This series represents environmental and climatic conditions over time. The final year (January to December 2016) was used as the test set.

Each series was preprocessed to ensure stationarity, applying differencing and seasonal adjustment techniques where necessary. The datasets were then split into training and testing sets, and model performance was evaluated to identify the best forecasting approach for each case.

2.2 Model specifications

To evaluate forecasting performance, two classes of models were considered:

- **SARIMA Models:** Seasonal Autoregressive Integrated Moving Average models were applied following the Box Jenkins methodology [3]. The general SARIMA model is denoted by SARIMA$(p, d, q)(P, D, Q)_s$, where p, d, and q are the non-seasonal autoregressive, differencing, and moving average orders, respectively; P, D, and Q are their seasonal counterparts; and s represents the seasonal period.

  The SARIMA model can be expressed as:

$$\Phi_P(L^s)\phi_p(L)(1-L)^d(1-L^s)^D y_t = \Theta_Q(L^s)\theta_q(L)\varepsilon_t,$$

  where:

  - L is the backshift operator: $L^k y_t = y_{t-k}$,
  - $\phi_p(L)$ and $\theta_q(L)$ are the non-seasonal AR and MA polynomials,
  - $\Phi_P(L^s)$ and $\Theta_Q(L^s)$ are the seasonal AR and MA polynomials,
  - ε_t is white noise.

The SARIMA model is particularly well-suited for time series that exhibit both trend and seasonality. It captures short-term autocorrelation structures through AR and MA terms and accounts for seasonal patterns via their seasonal counterparts.

Model identification was based on the analysis of autocorrelation (ACF) and partial autocorrelation (PACF) plots. For seasonal components (P and Q), attention was given to lags that are multiples of the seasonal period s. Optimal model orders were selected using the Akaike Information Criterion (AIC).

- **Artificial Neural Networks:** Feedforward neural networks with a single hidden layer were employed, adhering to the nonlinear autoregressive framework for univariate time series forecasting (Figure 2.1). In this structure, the current value of the time series is modeled as a nonlinear function of its past values (lags). Specifically, the model can be expressed as:

$$y_t = f(y_{t-1}, y_{t-2}, \dots, y_{t-p}) + \varepsilon_t,$$

where $f(\cdot)$ denotes a nonlinear mapping function approximated by the ANN, p is the number of lagged inputs (input neurons), and ε_t is the error term.

The architecture NNAR(p, k) consists of:

  - An input layer with p neurons representing past observations,
  - A hidden layer with k neurons using a nonlinear activation function (sigmoid by default),
  - An output layer with a single neuron with linear activation to produce real-valued forecasts.

The number of lagged inputs p and the number of hidden units k were determined through a trial-and-error approach, guided by performance on a validation set, and evaluated using the root mean square error (RMSE).

The models were trained using the backpropagation algorithm with a gradient descent optimizer. To prevent overfitting and improve generalization, early stopping was implemented based on validation performance. This means that during training, the model error in a validation set was continuously monitored. Once the error stopped decreasing and began to increase (indicating possible overfitting), training was stopped. This method allows the network to achieve optimal generalization by not overfitting the training data. ANN models were developed and evaluated using the `nnetar()` function from the `forecast` package in R, which can automatically select the appropriate input lag and fit a feedforward neural network with one hidden layer. By default, it averages forecasts from 20 independently trained networks to increase stability and reduce variance. This approach has been validated in various forecasting contexts, including recent studies such as [12], who successfully applied a NNAR(27,5) model to forecast COVID-19 cases in Italy.

To handle seasonality in the data, the NNAR framework extends to seasonal neural network autoregressive models, denoted as $\texttt{NNAR}(p, P, k)_s$, where p is the number of non-seasonal lags, P is the number of seasonal lags (i.e., multiples of the seasonal period s), and k is the number of neurons in the hidden layer. In this context, the model incorporates both recent and seasonal past observations as inputs:

$$y_t = f(y_{t-1}, \dots, y_{t-p}, y_{t-s}, \dots, y_{t-Ps}) + \varepsilon_t.$$

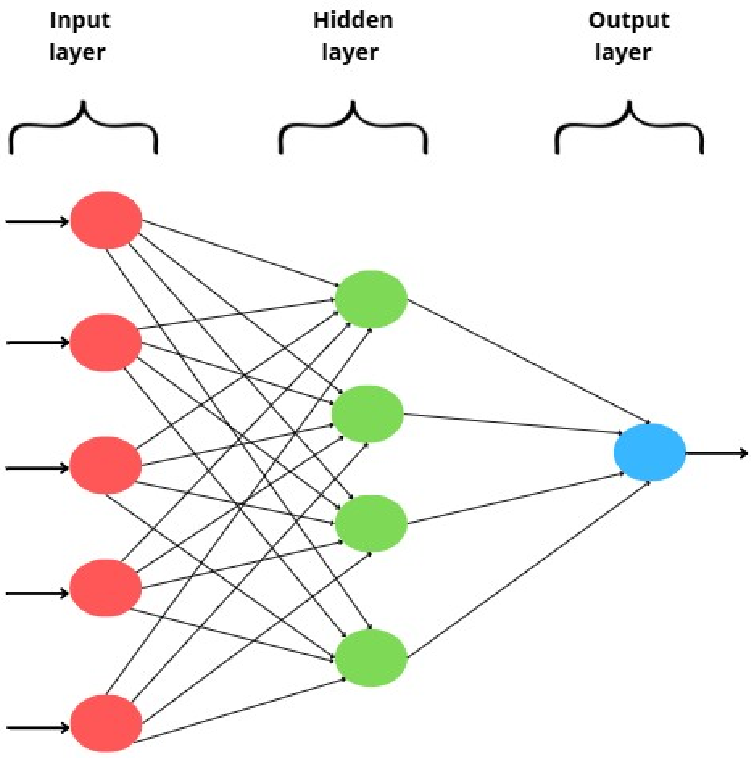

Figure 2.1: Architecture of the feedforward neural network model NNAR(5,4). The model includes 5 input neurons corresponding to lagged values, 4 neurons in the hidden layer with sigmoid activation, and 1 output neuron with linear activation.

This extension allows the network to effectively capture both short-term dependencies and recurring seasonal patterns in the data. The `nnetar()` function in R automatically detects and incorporates seasonality when the input time series exhibits periodic structure, making the seasonal NNAR model well-suited for monthly or quarterly data.

2.3 Performance metrics

To assess forecasting accuracy, we use the following standard performance metrics. Let y_t denote the actual value, $\hat{y}_t$ the forecasted value at time t, and n the number of forecasts in the test set.

- **Mean Error (ME):**

$$\text{ME} = \frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t).$$

- **Root Mean Square Error (RMSE):**

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2}.$$

- **Mean Absolute Error (MAE):**

$$\mathrm{MAE} = \frac{1}{n}\sum_{t=1}^{n} |y_t - \hat{y}_t| \,.$$

- **Mean Absolute Percentage Error (MAPE):**

$$\mathrm{MAPE} = \frac{100}{n}\sum_{t=1}^{n} \left| \frac{y_t - \hat{y}_t}{y_t} \right| .$$

These metrics were computed on the test datasets to evaluate out-of-sample performance. Comparisons between SARIMA and ANN forecasts were made both numerically and graphically to assess consistency across different types of data and forecasting horizons.

While confidence intervals are displayed for some forecasts, simulation plots are used in others to illustrate the range of possible trajectories generated by the NNAR model. This choice was intentional to highlight different aspects of model behavior interval uncertainty in some cases, and dynamic simulation capacity in others.

3 Results and discussion

3.1 Forecasting GDP series

We begin our analysis with the fitted neural network model applied to the training dataset. The model used is a NNAR(4,5), which corresponds to a feedforward neural network with 4 input neurons (representing the last 4 lagged observations), 5 neurons in the hidden layer, and 1 output neuron for forecasting.

The model was constructed as an average of 20 independently trained networks. Each individual network follows a 4-5-1 architecture and involves a total of 31 trainable weights (including biases). The residual variance $\hat{\sigma}^2$ is estimated to be 0.0063, indicating a relatively low level of unexplained variation in the training data, which suggests a good fit of the model. This model will now serve as the basis for forecasting and simulation.

The performance of both NNAR(4,5) and ARIMA(1,1,1) models was evaluated using the GDP time series. Table 3.1 presents the error metrics for both training and test sets. The NNAR model demonstrated superior performance across most accuracy measures. For the training set, the NNAR model yielded a significantly lower Root Mean Square Error (RMSE) of 0.0799 compared to 0.1289 for the ARIMA model. It also achieved lower Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE), indicating more precise and reliable forecasts. The Mean Error (ME) was nearly zero for NNAR, reflecting unbiased predictions, whereas the ARIMA model exhibited a slight positive bias.

On the test set, the NNAR model continued to outperform the ARIMA model. The RMSE and MAE for NNAR were 0.1246 and 0.1043, respectively, both of which are lower than those of ARIMA (0.2306 and 0.1736). Furthermore, the MAPE dropped to 1.95% for NNAR versus 3.23% for ARIMA, showing the robustness of the NNAR model in handling out-of-sample forecasts. The consistent reduction in all error metrics confirms the effectiveness of the NNAR model in modeling the nonlinear dynamics of Algeria's GDP series.

Figure 3.1 illustrates a graphical comparison between actual GDP values and the forecasts produced by the NNAR and ARIMA models. The actual series exhibits both upward and downward movements, reflecting the inherent nonlinear dynamics and structural changes

Table 3.1: Forecast accuracy measures for GDP time series

Model	Set	ME	RMSE	MAE	MAPE
NNAR	Training	-0.0003	0.0799	0.0605	1.6883
	Test	0.0711	0.1246	0.1043	1.9591
ARIMA	Training	0.0191	0.1289	0.0953	3.9419
	Test	0.1734	0.2306	0.1736	3.2351

within the economic data. The NNAR model captures the underlying trend of the series more accurately, particularly during the post-2020 period where the actual GDP shows a sharp increase. In contrast, the ARIMA model fails to adapt to the sudden rise, continuing a smooth downward trend that underestimates the actual GDP trajectory.

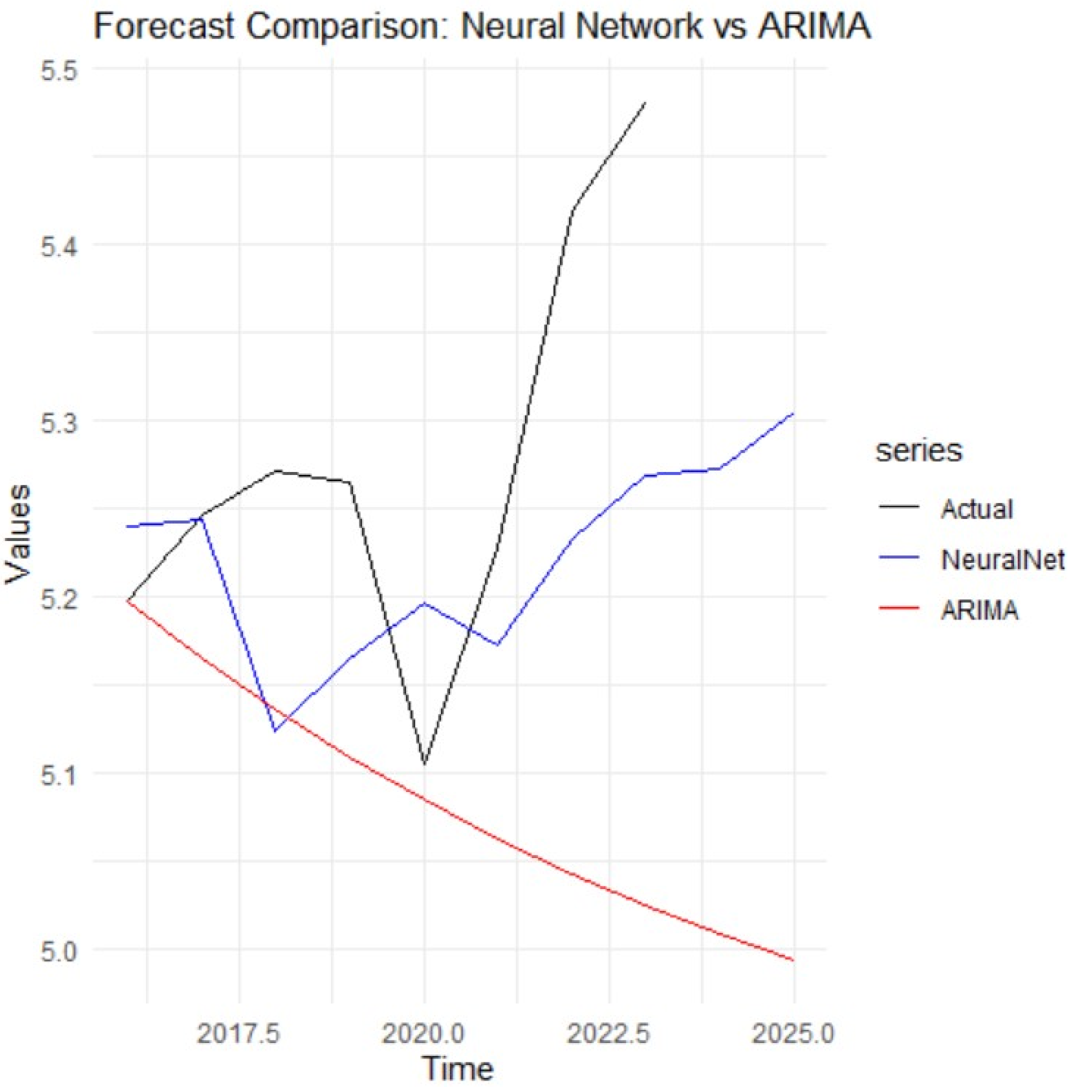

Figure 3.1: Forecast comparison of Algeria's GDP: NNAR vs ARIMA.

This visual evidence reinforces the results obtained from the error metrics. The consistent divergence of the ARIMA model from the real data after 2019 highlights its limitations in capturing nonlinearity, making it less suitable for GDP series characterized by structural breaks or regime shifts.

Figure 3.2 shows residual analysis of the ARIMA(1,1,1) model. The residuals were first plotted along with their ACF to visually check for any remaining autocorrelation. The ACF plots of residuals fall within the 95% confidence intervals, indicating no significant autocorrelation. Furthermore, we performed the Box Ljung test to formally assess the independence of residuals, and the Jarque Bera test to examine their normality. The Box Ljung test gives p-value = 0.7009, and for the Jarque Bera test, the p-value was 0.0629. Thus, at the 5% significance level, we fail to reject the null hypothesis of normality, but at the 10% level, we reject it.

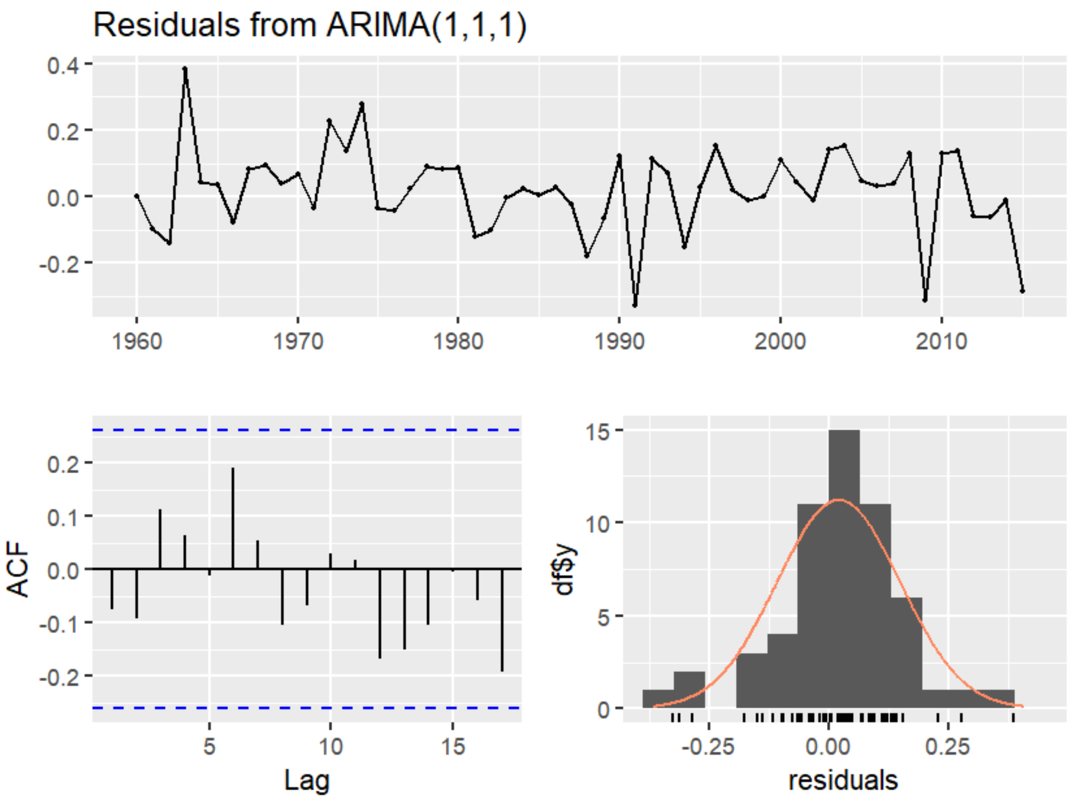

Figure 3.2: SARIMA model diagnostics for Algeria's GDP series. (Top) Residuals plot showing no apparent patterns. (Middle) ACF of residuals with 95% confidence bands. (Bottom) Histogram of residuals with normal distribution overlay.

To explore potential future scenarios, we simulated nine possible future paths using the fitted NNAR(4,5) model based on the entire dataset. As shown in Figure 3.3, the black curve depicts the historical evolution of the logarithm of the observed time series, while the colored lines represent the simulated trajectories.

The superior performance of the NNAR model may be attributed to its ability to capture nonlinear patterns and structural breaks in the Algerian economy, especially after 2020, when oil price shocks and fiscal reforms significantly altered the dynamics of GDP growth. Unlike the linear SARIMA model, the NNAR approach adapts more flexibly to these abrupt regime shifts.

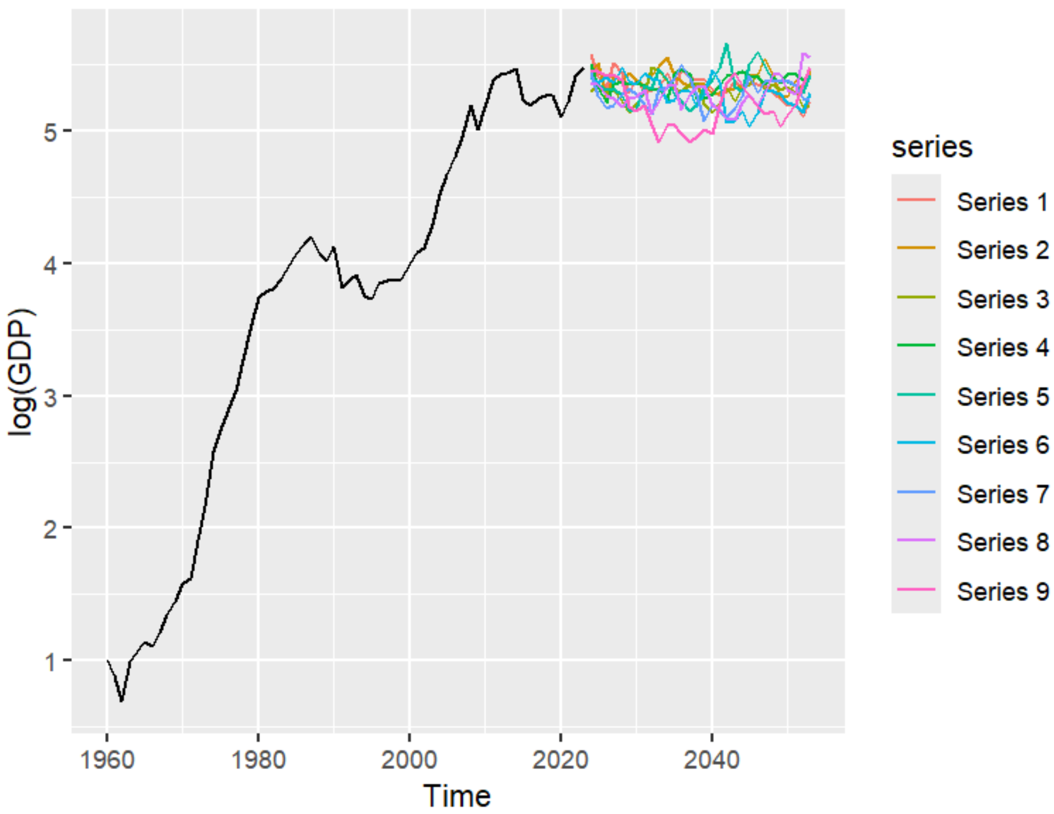

Figure 3.3: Forecasting simulations from the fitted NNAR(4,5) model for Algeria's GDP. The black line shows historical log-transformed GDP data, while colored lines represent nine simulated future trajectories with prediction intervals.

3.2 Temperature forecasting results

The selected model for the temperature time series is a NNAR(11,1,10)[12], comprising an ensemble average of 20 neural networks. Each network follows a 12-10-1 architecture, with 12 input nodes (11 lagged values plus one seasonal lag at period 12, capturing both short-term and seasonal dependencies), 10 hidden nodes, and one output node. Each network contains 141 trainable weights. The relatively low estimated noise variance ($\hat{\sigma}^2 = 0.0006$) suggests a strong in sample fit.

Table 3.2 summarizes the performance metrics of the NNAR(11,1,10)[12] and SARIMA(1,0,0)-(0,1,1)[12] models. The NNAR model outperforms SARIMA on both training and test sets. It yields lower errors across all metrics, indicating better fit and more accurate forecasts. This highlights the superior ability of the NNAR model to capture the underlying patterns of the series.

Figure 3.4 presents the forecast comparison between the NNAR model and the SARIMA model for the temperature time series. Both models accurately capture the seasonal behavior and the overall trend of the data. The Neural Network model follows the actual observations slightly more closely, especially during the phases of rapid temperature increase and decrease. SARIMA also performs well but shows small deviations, particularly at the beginning.

Figure 3.5 presents the residual analysis of the SARIMA model. The Box Ljung test yields a

Table 3.2: Performance metrics for temperature forecasting using NNAR and SARIMA models

Model	Set	ME	RMSE	MAE	MAPE
NNAR	Training	0.0001	0.0254	0.0192	0.6590
	Test	0.0285	0.0436	0.0300	1.0579
SARIMA	Training	0.0074	0.0393	0.0281	0.9671
	Test	0.0299	0.0463	0.0320	1.1220

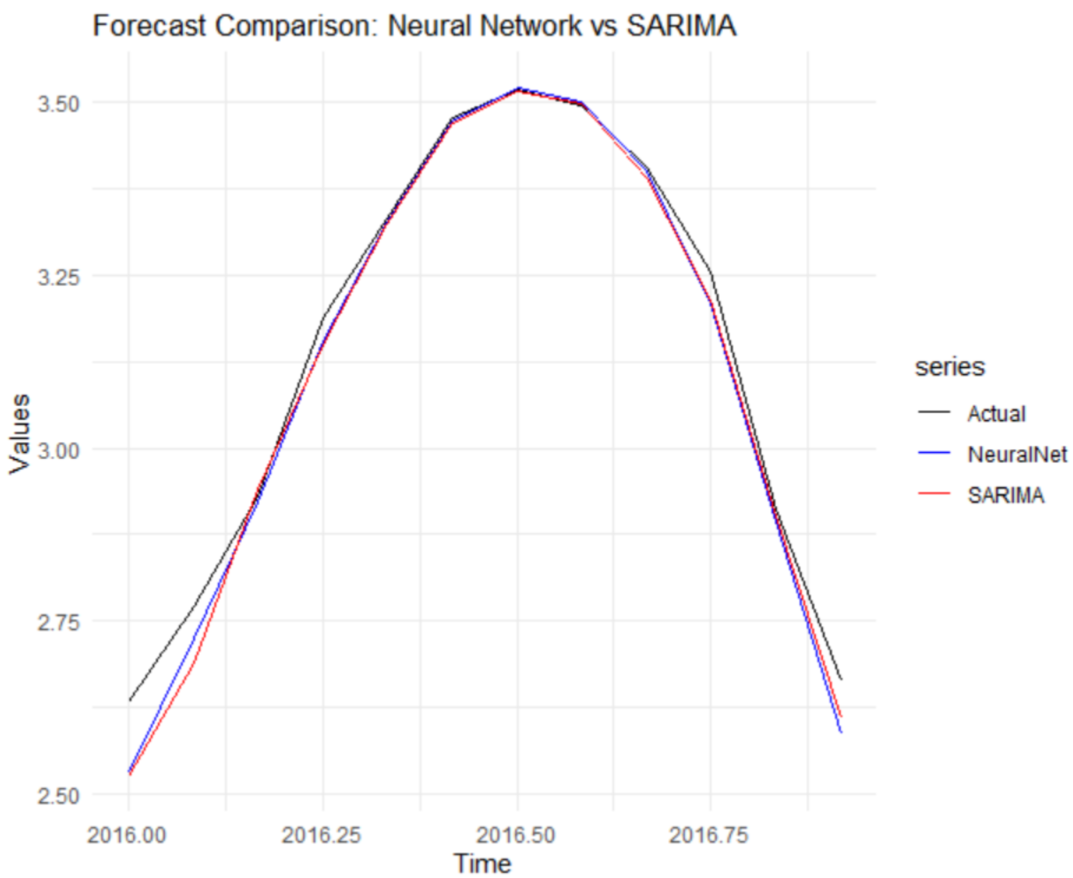

Figure 3.4: Forecast comparison between NNAR and SARIMA models for Algeria's monthly average temperature.

p-value of 0.9324, suggesting no significant autocorrelation in the residuals. However, the Jarque Bera test indicates strong evidence of non-normality (p-value $< 2.2 \times 10^{-16}$), supporting the presence of nonlinearity in the data.

Figure 3.6 displays nine simulated trajectories generated from the NNAR(11,1,10)[12] model fitted to the full dataset.

The stronger NNAR performance for temperature data likely reflects underlying climatic nonlinearities influenced by regional and global drivers, which are better captured by neural networks than by linear models.

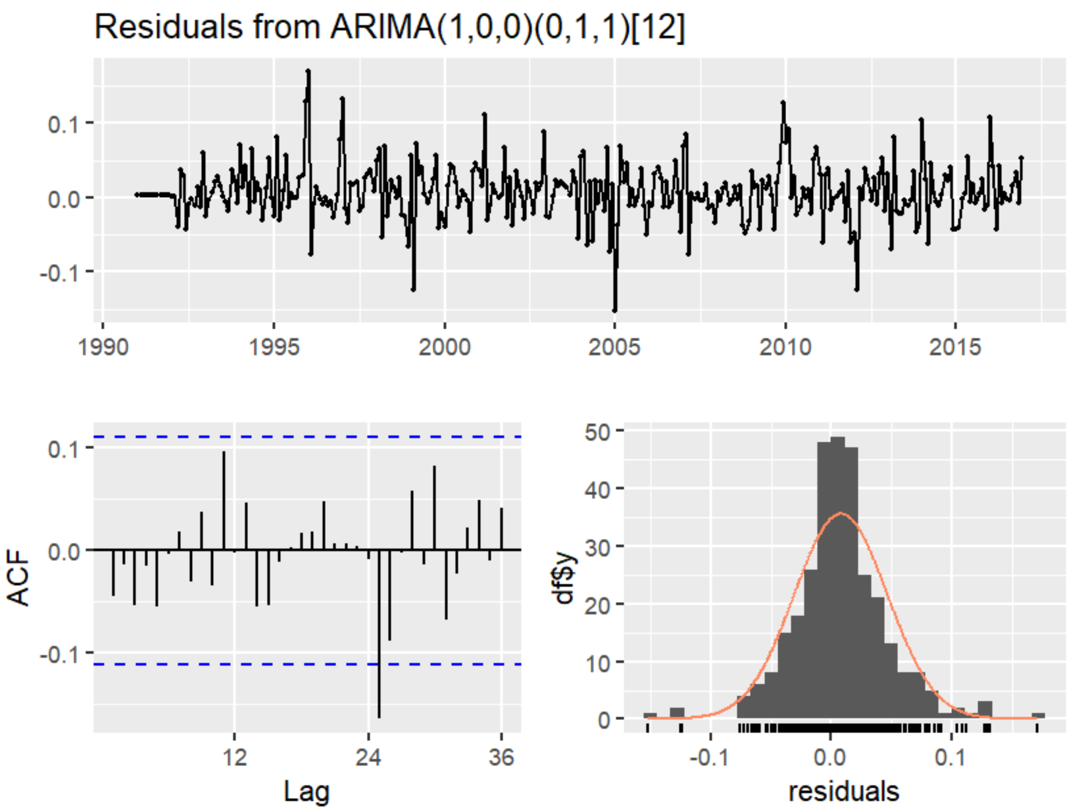

Figure 3.5: SARIMA model diagnostics for Algeria's temperature series. (Top) Residuals plot. (Middle) ACF of residuals with 95% confidence bands. (Bottom) Histogram of residuals with normal distribution overlay.

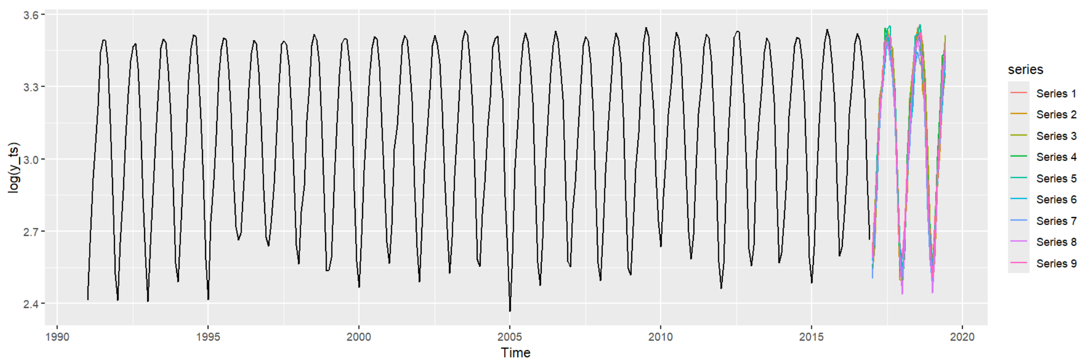

Figure 3.6: Forecasting simulations from the fitted NNAR(11,1,10)[12] model for Algeria's temperature. The black line shows historical temperature data, while colored lines represent nine simulated future trajectories with prediction intervals.

3.3 Forecasting USD/DZD exchange rate

The third application focuses on modeling the daily USD/DZD exchange rate from January 1, 2022 to January 1, 2024 (approximately 5 observations per week, excluding weekends). The selected model was a NNAR(27,1,5)[5], indicating the use of 27 non-seasonal lags and 1 seasonal lag (with a seasonal period of 5, corresponding to weekly seasonality), and a hidden layer comprising 5 neurons. The model was trained as the average of 20 independently initialized feedforward networks, each having a 27-5-1 structure and a total of 146 weights.

The estimated noise variance was $\hat{\sigma}^2 = 3.31 \times 10^{-6}$, suggesting a tight fit to the training data. The relatively large number of input lags allows the model to capture the potentially long-memory behavior and complex nonlinear dynamics typical of exchange rate series. The inclusion of weekly seasonality is particularly relevant in financial time series, where cyclical effects may appear due to trading behavior and institutional activity concentrated during business days.

Table 3.3 summarizes the obtained results. The NNAR model clearly outperforms the SARIMA(2,1,0)(1,0,1)[5] model on both training and test sets: Training performance is better for NNAR, with significantly lower RMSE and MAPE. The lower ME values suggest negligible bias for both models, but the smaller RMSE and MAE for NNAR on the test set confirm better generalization.

Table 3.3: Performance comparison between NNAR and ARIMA models on the USD/DZD exchange rate series.

Model	**Set**	**ME**	**RMSE**	**MAE**	**MAPE**
NNAR	Training	-2.20×10^{-6}	0.0018	0.0013	0.0270
	Test	8.76×10^{-5}	0.0057	0.0048	0.0995
SARIMA	Training	-4.82×10^{-6}	0.0026	0.0019	0.0396
	Test	-1.59×10^{-2}	0.0195	0.0164	0.3353

Figure 3.7 displays the forecast comparison between the actual USD/DZD values and the forecasts obtained from the NNAR and SARIMA models. It can be observed that the NNAR model follows the downward trend of the actual series more closely than the SARIMA model, which remains almost flat and fails to capture the dynamic changes. This graphical evidence supports the numerical results, confirming the superior forecasting performance of the NNAR model for this exchange rate series.

The residual plots do not display any visible patterns, suggesting no major model misspecification (Figure 3.8). The p-value of the Box Ljung test is 0.6048, confirming the null hypothesis of independence. Therefore, the fitted SARIMA model appears to be adequate for the data under study. As financial data, the p-value of the Jarque Bera test is $< 2.2 \times 10^{-16}$, confirming significant deviation from normality.

Figure 3.9 presents forecasts with prediction intervals generated by the NNAR(27,1,5)[5] model fitted to the complete dataset.

The poor performance of SARIMA for the USD/DZD exchange rate likely reflects the nonlinear and heteroskedastic nature of financial data. While models such as GARCH, RNN, or LSTM could better accommodate these dynamics, the present study focuses on comparing NNAR and SARIMA under a common framework, leaving the exploration of volatility-based or deep learning extensions for future work.

4 Conclusion

This study evaluated the forecasting performance of the Neural Network Autoregressive (NNAR) model in comparison with the classical SARIMA model. Using real-world Algerian datasets, the NNAR model consistently delivered more accurate forecasts, especially in contexts involving nonlinear patterns and complex dynamics.

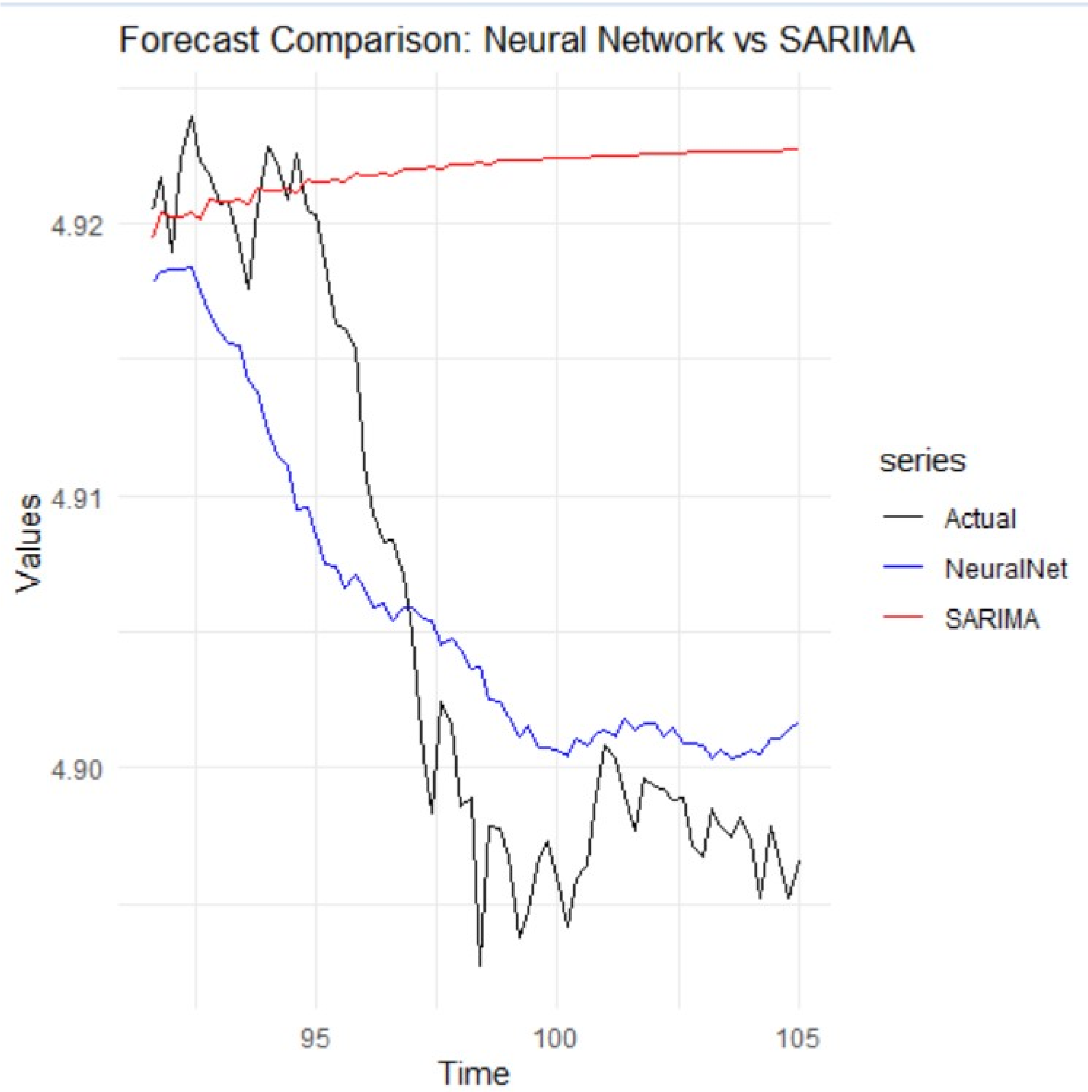

Figure 3.7: Forecast comparison between NNAR and SARIMA models for the USD/DZD exchange rate.

The results confirm the flexibility and robustness of NNAR in modeling time series data, making it a strong alternative to SARIMA. Its ability to adapt to nonlinearity and irregular structures enhances its predictive performance, particularly for financial and economic data.

Future research will aim to extend this work by exploring more advanced neural network architectures such as RNNs and LSTM networks. In addition, hybrid models that combine the strengths of neural networks and traditional models may offer even greater accuracy and reliability for forecasting.

Declarations

Availability of data and materials

The datasets used in this study are publicly available from the sources cited in Section 2.1.

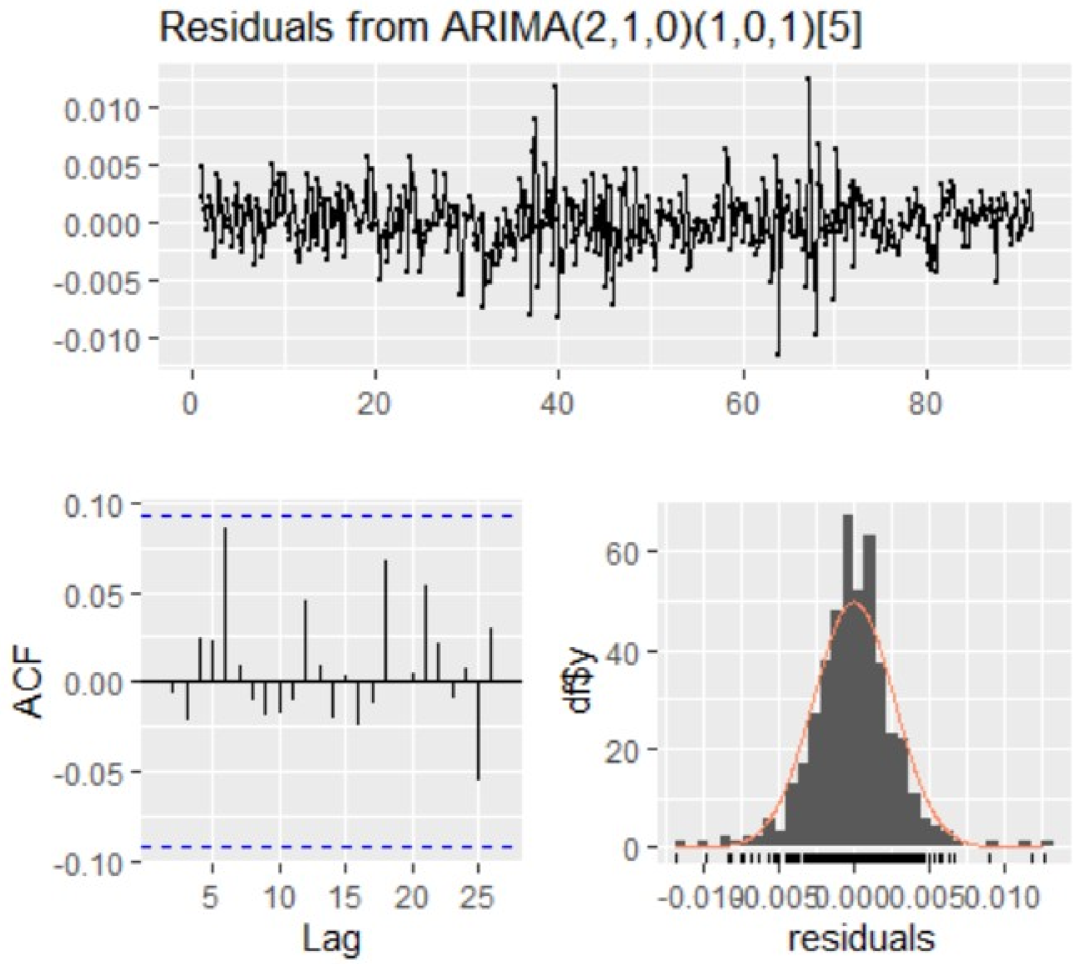

Figure 3.8: SARIMA model diagnostics for USD/DZD exchange rate data. (Top) Residuals plot. (Middle) ACF of residuals with 95% confidence bands. (Bottom) Histogram of residuals with normal distribution overlay.

Funding

Not applicable.

Authors' contributions

K. Djaber: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing original draft, Visualization. M. Merzougui: Conceptualization, Methodology, Resources, Writing review & editing, Supervision, Project administration.

Conflict of interest

The authors have no conflicts of interest to declare.

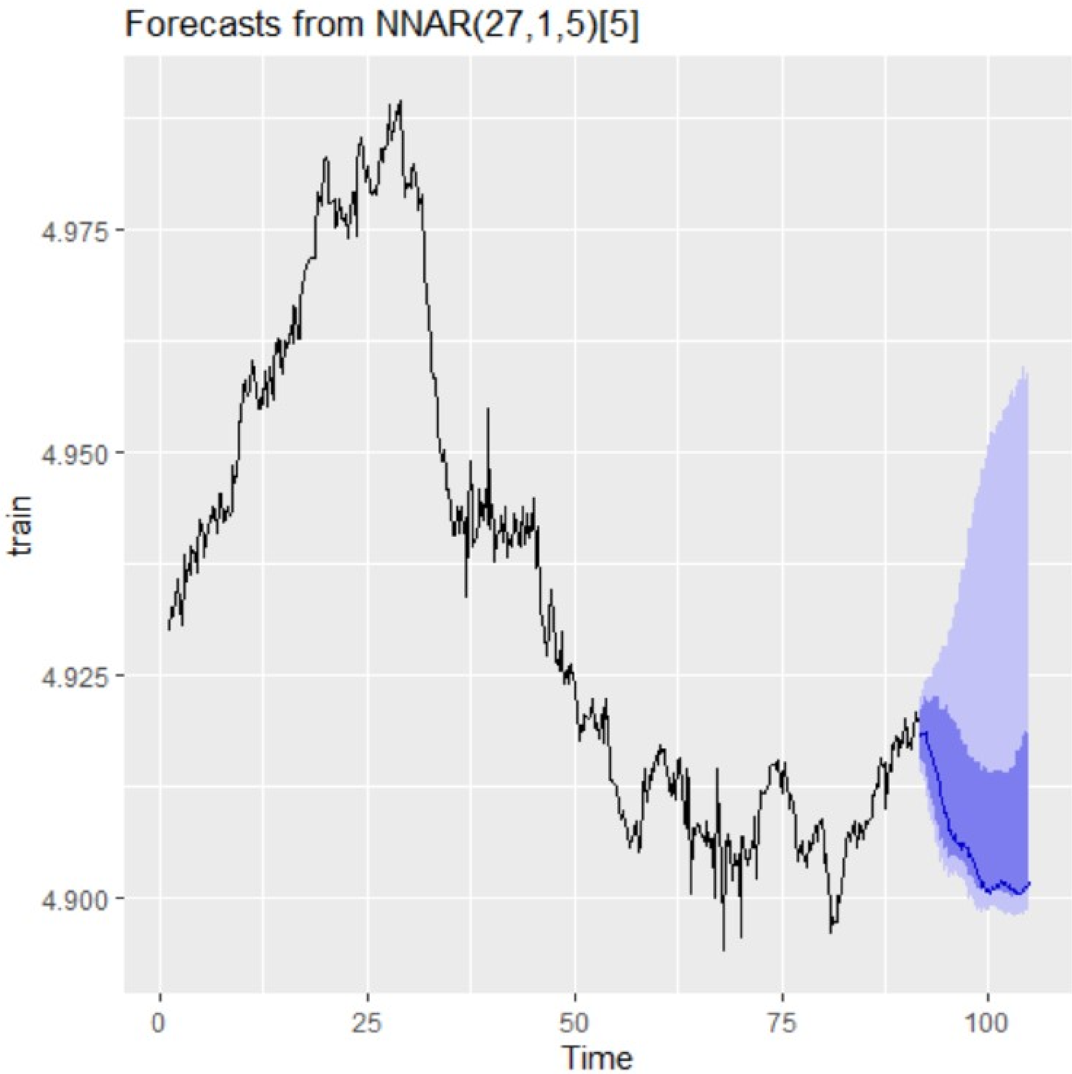

Figure 3.9: Forecasts from the fitted NNAR(27,1,5)[5] model for USD/DZD exchange rate. The black line shows historical exchange rate data, while the blue line represents point forecasts with 80% and 95% prediction intervals shown in shaded bands.

Acknowledgements

The authors acknowledge the valuable comments of the reviewers and the institutional support of Badji Mokhtar Annaba University.

References

[1] A. Benlaria and L. Boubekeur, *Forecasting exchange rates using artificial neural networks*, Majallat al-Basha'ir al-Iqtisadiyya, **7**(2) (2021), 804–814.

[2] C. M. Bishop, *Neural networks for pattern recognition*, Oxford University Press, 1995.

[3] G. E. P. Box, G. M. Jenkins and G. C. Reinsel, *Time series analysis: Forecasting and control*, Prentice-Hall, 1994.

[4] I. Elbatal, M. Sarwar, F. Jamal, M. Daniyal, Z. Hussain and A. Ben Ghorbal, *Modeling on gross domestic product annual growth rate data by using time series, machine learning, and probability models*, Journal of Radiation Research and Applied Sciences, **18** (2025), 101481. DOI

[5] I. Goodfellow, Y. Bengio and A. Courville, *Deep learning*, MIT Press, 2016. http://www.deeplearningbook.org

[6] L. Hussain, B. Ghufran and A. Ditta, *Forecasting inflation, exchange rate, and GDP using ANN and ARIMA models: Evidence from Pakistan*, Sustainable Business and Society in Emerging Economies, **4**(1) (2022).

[7] R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and practice* (2nd ed.), OTexts, 2018. https://otexts.com/fpp2/

[8] S. Lakhal, *Forecasting the EUR/USD exchange rate using ARIMA and machine learning models*, Data and Metadata, **3** (2024), 368. DOI

[9] T. M. Mitchell, *Machine learning*, McGraw-Hill, 1997.

[10] R. M. Nezzar, N. Farah, M. T. Khadir and L. Chouireb, *Mid-long term load forecasting using multi-model artificial neural networks*, International Journal on Electrical Engineering and Informatics, **8**(2) (2016), 317–329.

[11] T. Raza Fraz, *Modeling and forecasting of rainfall time series: A case study for Pakistan*, International Journal of Economic and Environmental Geology, **13**(1) (2022), 37–41.

[12] L. Saliaj and E. Nissi, *Artificial neural networks for COVID-19 time series forecasting*, Open Journal of Statistics, **12** (2022), 277–290. DOI

[13] H. Şengüler and B. Kara, *Forecasting the inflation for budget forecasters: An analysis of ANN model performance in Türkiye*, Journal of Research in Economics, Politics and Finance, **10**(1) (2025), 58–91.

Appendix

This appendix provides supplementary material supporting the empirical analysis, including visual diagnostics of non-stationarity for the three Algerian time series, forecast accuracy comparison through Theil's U statistic, and the architecture of the fitted NNAR model.

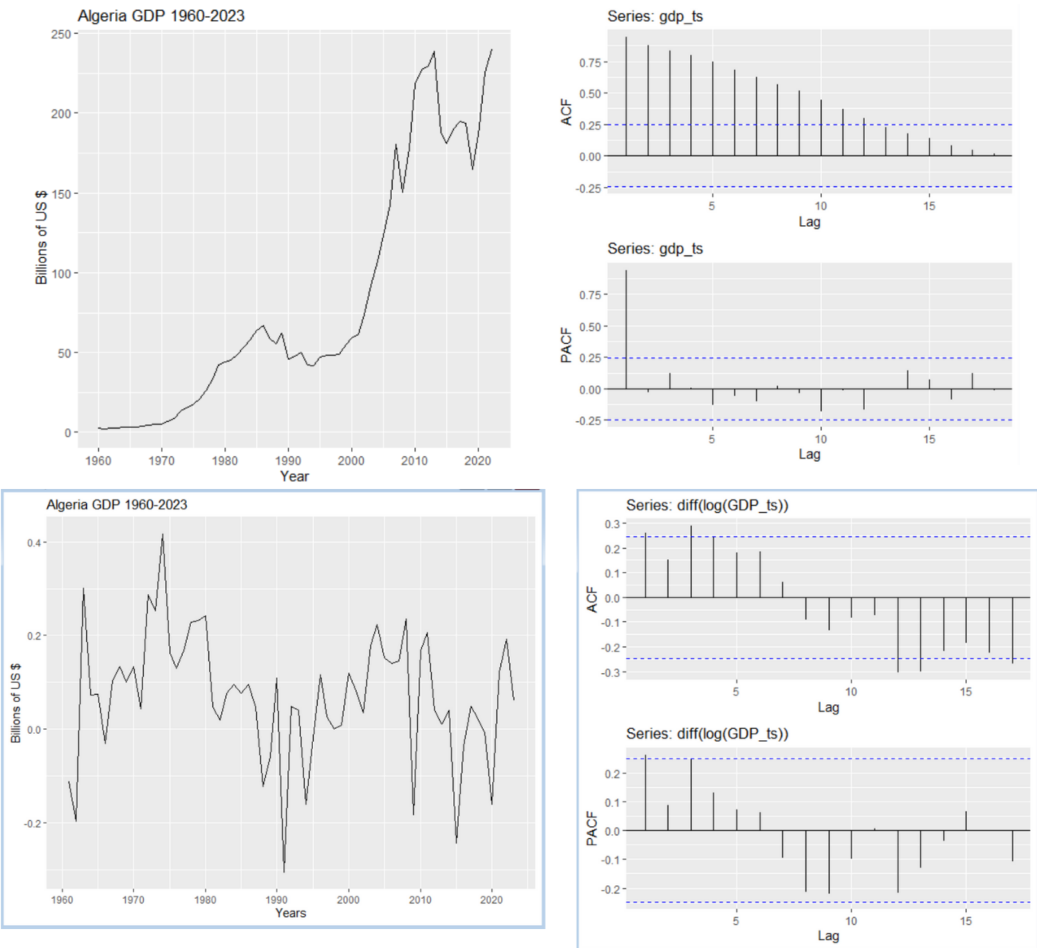

Figure .1: Time series analysis of Algeria's GDP from 1960 to 2023. (Top row) Raw GDP data and log-transformed series. (Middle row) First-differenced log GDP series. (Bottom row) ACF and PACF plots of the differenced series.

Table .1: Theil's U statistic for the three Algerian time series

Series	NNAR	SARIMA	Interpretation
GDP (log)	0.9649	2.2142	NNAR yields a lower Theil's U, indicating more accurate and consistent forecasts than SARIMA.
Exchange rate (USD/DZD, log)	2.9369	9.7234	Both models exhibit limited accuracy, but NNAR still performs better, with substantially lower forecast error proportionality.
Temperature (log)	0.1779	0.2107	Both models perform well, yet NNAR slightly outperforms SARIMA, showing improved short-term predictive accuracy.

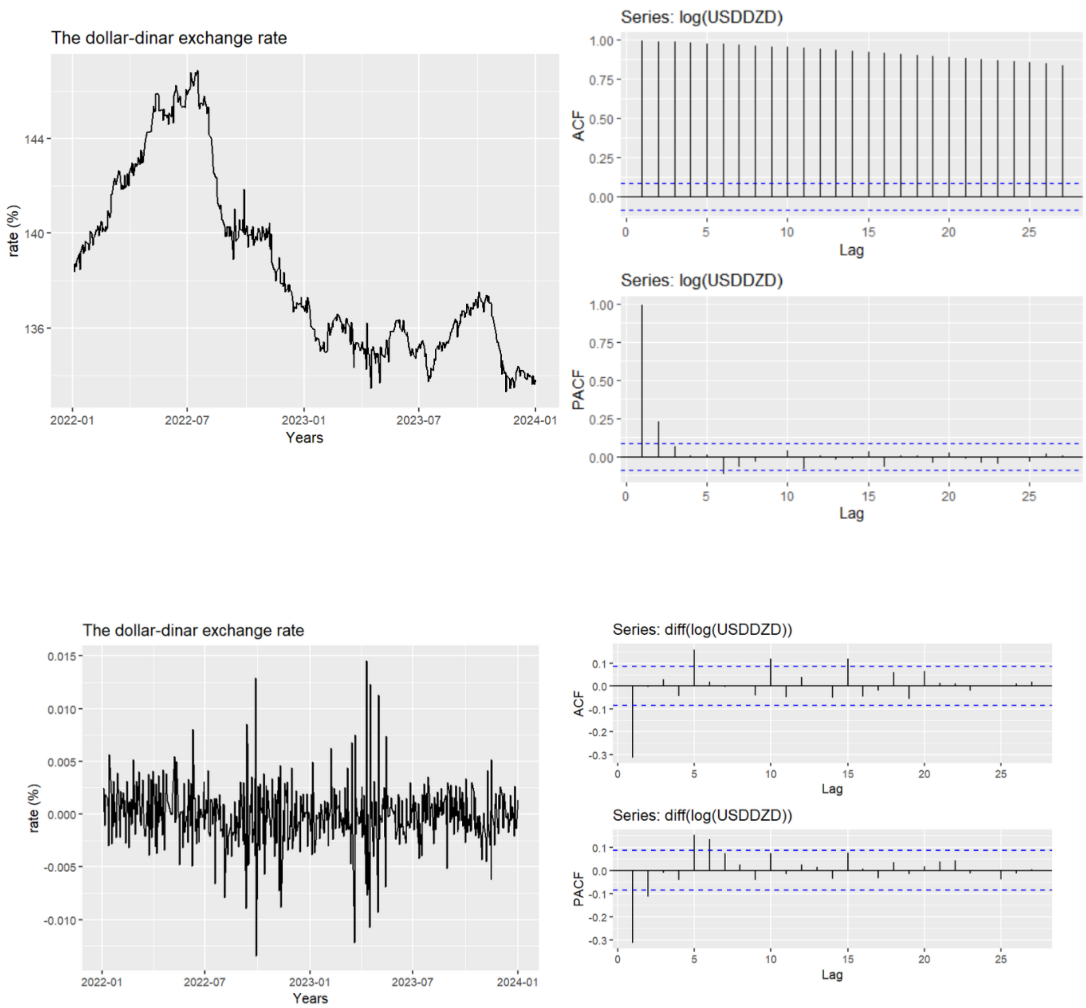

Figure .2: Time series analysis of the USD/DZD exchange rate from 01/01/2022 to 01/01/2024. (Top row) Raw exchange rate data and log-transformed series. (Middle row) First-differenced log exchange rate series. (Bottom row) ACF and PACF plots of the differenced series.

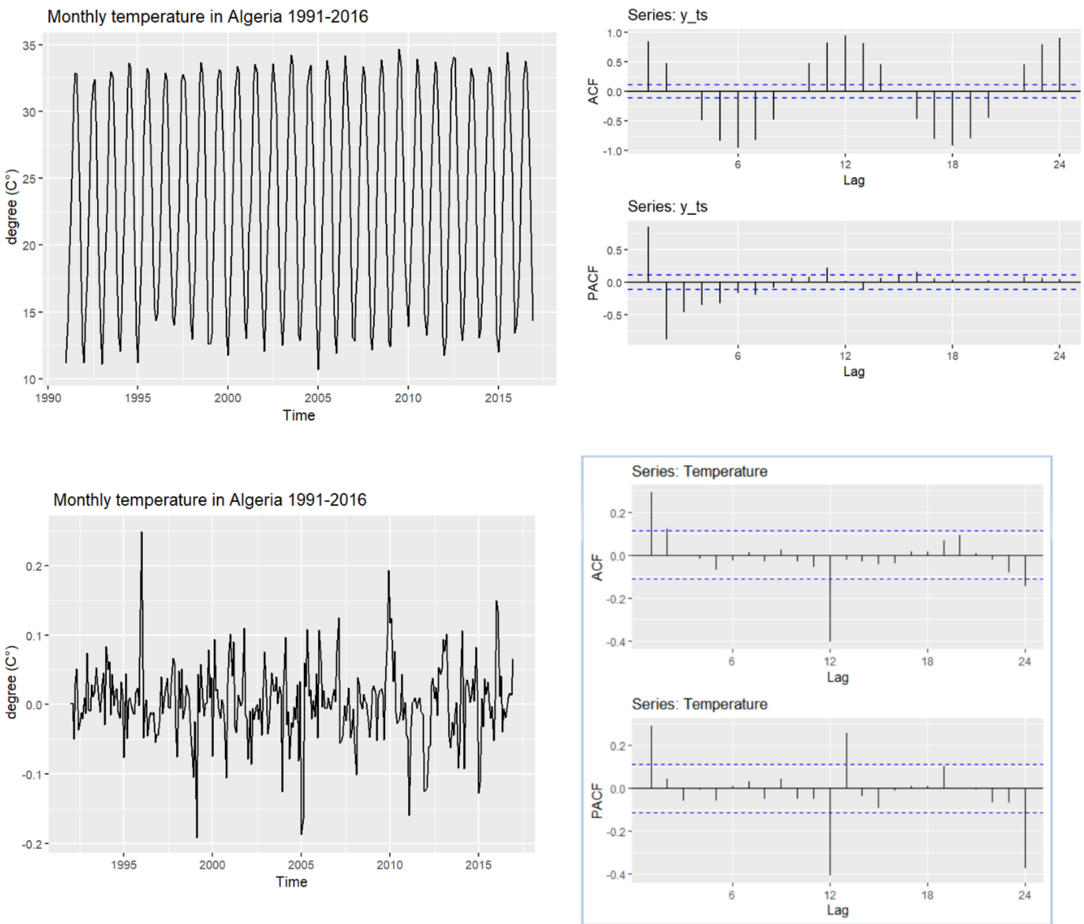

Figure .3: Time series analysis of monthly temperature in Algeria from 01/1991 to 12/2016. (Top row) Raw temperature data showing seasonal patterns. (Middle row) Seasonally differenced temperature series. (Bottom row) ACF and PACF plots of the seasonally differenced series.

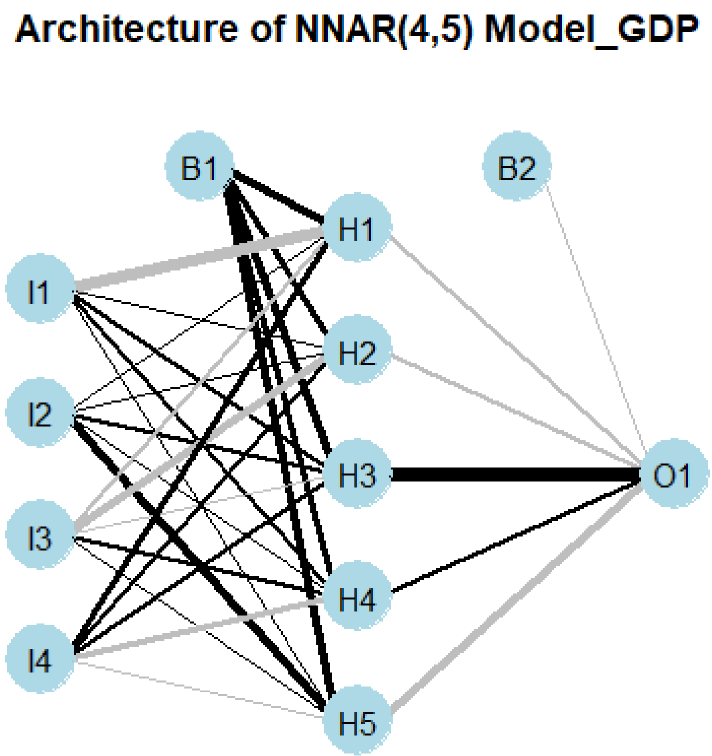

Figure .4: Architecture of the NNAR(4,5) model fitted to Algeria's GDP. The model includes four input neurons corresponding to four lagged values of the log-transformed GDP, five neurons in the hidden layer, and one output neuron. The thickness of the connections reflects the relative weight magnitudes, as visualized by the NeuralNetTools package in R.