ISSN : 3105-0794

الهيئة الليبية للبحث العلمي

Libyan Authority For Scientific Research

Libyan International Journal of Natural Sciences
2025
LIJNS
Peer-Reviewed Journal

(LIJNS)

Libyan International Journal of Natural Sciences

VOLUME 2 ISSUE 1 MAY 2026

Bi-Annual, Peer- Reviewed, and Open Accessed e-Journal

DOLJ
Directory of Online Libyan Journals

turnitin

Google Scholar

Puplished by

MILJS@aonsrt.ly

Libyan International Journal of Natural Sciences (LIJNS)
Vol.2 Issue No 1 May 2026

Forecasting Monthly Air Temperature in Northeastern Libya Using Regularised Regression Models.

Jamila Omar Aldrogi
Department of Statistics, Faculty of Science, University of Benghazi, Benghazi, Libya.

ARTICLE INFORMATION

Article history:
Received 18 February 2026
Accepted 27 March 2026
Published 14 April 2026

Keywords:
Time series forecasting; Temperature prediction; Regularized regression; Fourier seasonality; NASA POWER; Climate modeling; Lag-based models

A B S T R A C T

This study addresses the need for accurate temperature forecasting in semi-arid Mediterranean regions to support climate adaptation, energy management. The objective is to develop a structured and computationally efficient forecasting framework for monthly air temperature in northeastern Libya (32.1167°N, 20.0667°E) using NASA POWER data spanning 2000–2024. To reduce high-frequency variability and enhance signal stability, daily meteorological observations were aggregated into monthly averages. Forecasting was formulated as a supervised learning problem using engineered lag features of the target variable alongside exogenous atmospheric drivers, including relative humidity (RH2M), wind speed (WS2M), and surface solar radiation (ALLSKY_SFC_SW_DWN). Seasonal dynamics were explicitly modeled using Fourier terms to capture annual periodicity. Three regularised regression models—Ridge, Lasso, and Elastic Net—were evaluated under a strict chronological framework comprising training (2000–2018), validation (2019–2021), and independent testing (2022–2024), with an expanding-window walk-forward strategy for one- and three-month-ahead forecasting. The results demonstrate that the Lasso-based model achieves the best overall performance, with a Root Mean Square Error (RMSE) of approximately 1.02°C for one-month-ahead forecasts and comparable accuracy for three-month-ahead predictions. While the seasonal naïve model remains competitive due to strong annual periodicity, the proposed framework consistently outperforms naïve benchmarks. Diebold–Mariano statistical tests confirm that improvements over the naïve persistence model are statistically significant ($p < 0.001$), although differences relative to the seasonal naïve model are not statistically significant. Regularised regression with structured lag features and explicit seasonal representation provides a robust, interpretable, and computationally efficient alternative to more complex nonlinear forecasting models for regional climate applications.

Introduction

Time series forecasting is of great importance in environmental modeling, energy management, and the mitigation of climate risks. Reliable temperature prediction can be used to support electricity demand planning, agriculture decision making, water resource management, and long-term stability of infrastructures in semi-arid Mediterranean climatic conditions. Strong annual seasonality, persistence, and meteorological dependence of atmospheric variables such as humidity, wind speed, and solar radiation are typical features of air-temperature series every month. Such structural aspects require forecasting models that may be used to represent the temporal dependency and exogenous climatic impacts collectively.

Even though naive and seasonal naive models remain the standard benchmark baseline models due to their simplicity and interpretability, the degree of their power is attributed to their limitation by definition. It is important to know that naïve seasonal forecasting can be competitive in regions with significant annual cycles; these models do not explicitly model dynamic interactions of temperature and meteorological drivers. It implies that they may fail to record small changes in normal season patterns when there are changes in the atmosphere.

Recent further developments on regression-based time series forecasting have demonstrated that structured lag representations may be employed to reduce temporal prediction to a supervised learning problem

* Corresponding author: *E*-mail addresses: jameelah.omar.edu.ly.

effectively. These models, together with regularization techniques, provide a consistent estimation of the coefficients in the case of multicollinearity between lagged predictors. Ridge, Lasso, and Elastic net regression are variants of classical linear regression that add penalty terms to the model that improve the generalization performance and minimize overfitting in moderately high-dimensional lag spaces.

It is not always that purely autoregressive models are necessarily sufficient to model strongly seasonally dependent climate series without any explicit representation of seasonal dynamics. Fourier seasonal terms are an elasticity, but not a very parameterized way of getting periodic structure in a model, coupled with lagged meteorological drivers.

In this paper, the superior regression model for the month-ahead air-temperature prediction in northeast Libya (32.1167 o N, 20.0667 o E) in NASA POWER climate data from 2000 to 2024 is constructed. The solution suggested encompasses:

Structured representations of temperature and exogenous drivers in lag-based form.

Elastic Net regularised regression models Ridge, Lasso, Elastic Net regularised regression models.

Fourier seasonality elements.

Critical chromological training-validation-testing design.

Increasing-window (walk-forward) out-of-sample testing.

Multi-horizon (one and three months) forecasting.

Diebold-Mario statistical comparison test.

Research Objective and Question

To address the limitations of traditional and baseline forecasting approaches in strongly seasonal climate conditions, this study aims to evaluate whether structured regularised regression models can provide statistically and practically significant improvements in temperature prediction accuracy. Accordingly, the primary research question of this study is:

Can regularized regression models augmented with structured lag features and Fourier-based seasonal components significantly outperform naïve and seasonal naïve benchmarks in monthly temperature forecasting under semi-arid Mediterranean climatic conditions?

To further refine this investigation, the study also examines:

Whether incorporating exogenous meteorological variables improves predictive performance.

Whether model performance remains stable across multiple forecasting horizons (h = 1 and h = 3 months).

Whether observed improvements are statistically significant based on the Diebold–Mariano test.

2. Related Work

The time series forecasting approaches may be broadly categorized into statistical autoregressive models, regression frameworks involving exogenous inputs, and nonlinear machine learning methods. The ARIMA and SARIMA models are classified as classical autoregressive integrated moving models of environmental and climatological forecasting because they stand on a strong theoretical base and can be used to characterize the structures of dependence over time. Of strong importance is the use of seasonal naive models, which can be applied as strong baselines in temperature forecast applications, as there is a high level of periodicity in annual variations. However, not all dynamic relationships between temperature and such meteorological parameters as relative humidity, wind speed, and solar radiation can be seasonally or autoregressively modeled.

To counter these limitations, lagged exogenous regression procedures are increasingly becoming popular. The linear modeling model can be built to develop forecasting through the creation of a structured lag of the explanatory and target variables without compromising the interpretation of the latter.

Ridge, Lasso, and Elastic Net regression are regularization methods that overcome these problems by having the magnitude of the coefficients penalized, thereby stabilizing the estimation of the parameters and improving out-of-sample prediction. Ridge regression minimizes L 2 penalization, Lasso permits the sparse selection of the variables by L 1 penalization, and Elastic Net strikes a balance between the two penalizations. It is these properties that make regularised regression particularly suited to address time series problems that include correlated lags as forms of problems.

Although the predictive capacity of deep learning architectures such as recurrent neural networks and long short-term memory (LSTM) models has been demonstrated to be effective on large-scale forecasting tasks, this costly and interpretable architecture may limit their applicability in the context of operational climate. In comparison, regularised regression models offer a good tradeoff between accuracy, transparency, and computational efficiency.

The present research contributes to this work by performing a systematic evaluation of regularised regression models augmented with Fourier seasonal terms with a stringent walk-forward forecasting design. The paper provides an empirically grounded assessment of the best-superior regression models to forecast monthly change in temperature in a highly seasonal Mediterranean climate through a comparison of performance across various levels as well as testing the statistical significance through the Diebold-Mariano test.

3. Materials and Methods

3.1 Study Area and Data Source

The meteorological data of the northeastern part of Libya (32.1167 o N, 20.0667 o E) was retrieved on the NASA Prediction of Worldwide Energy Resources (POWER) site.

The covered data is on a daily basis, starting January 1, 2000, to December 31, 2024.

In order to decrease high-frequency variability and improve signal stability, daily observations were pooled into monthly means, which created a continuous time series of 300 monthly observations.

The target variable is:

T2M: 2-meter air temperature (°C)

Exogenous meteorological drivers are:

RH2M: refer to the relative humidity at 2 meters (%)

WS2M: Wind speed at 2 meters (m/s)

ALLSKY_SFC_SW_DWN: refre Surface shortwave solar radiation (W/m²)

Precipitation (PRECTOT) was initially considered but excluded after preliminary validation indicated limited incremental contribution to monthly temperature predictability.

The data was sorted in time order and was not randomly shuffled as a time-dependent sequence.

Figure 1 presents the monthly 2-meter air temperature (T2M) series for the period 2000–2024. The series exhibits pronounced annual oscillations and moderate interannual variability characteristic of semi-arid Mediterranean climates.

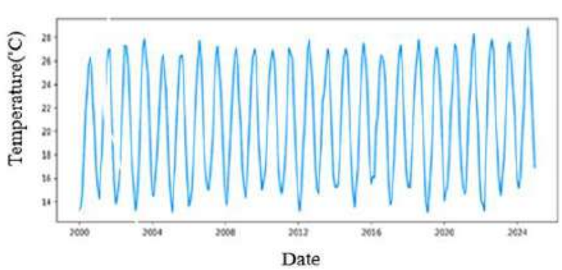

Fig.1: (T2M) monthly air temperature in northeastern Libya (2000–2024).

In order to investigate structural elements of the series further, STL decomposition was used to extract trend, seasonal, and residual elements (Figure 2). Strong periodicity is confirmed by the fact that the variability is dictated by a seasonal component.

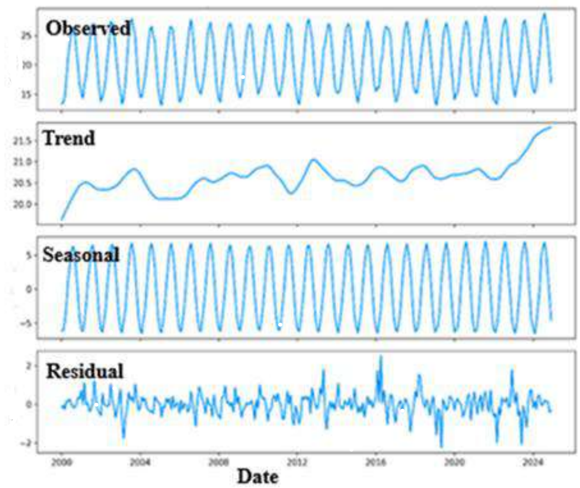

Fig. 2: Decomposition of monthly temperature series (observed, trend, seasonal, and residual components).

Fig. 3: Demonstrates the distribution of temperature by calendar month, which also demonstrates consistent behaviour on a year-by-year basis.

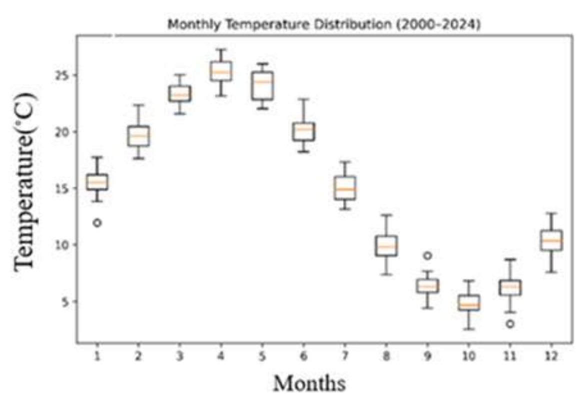

Fig.3: Monthly temperature by calendar month distribution (2000–2024).

3.2 Data Preprocessing

3.2.1 Missing Value Treatment

Where there were missing observations, time-consistent linear interpolation was employed within the temporal observations adjacent to the missing ones only. Imputation did not include any future information that would cause leakage.

3.2.2 Standardization
All predictor variables were standardized using training-set statistics only:

$$z_t = \frac{x_t - \mu\ \text{train}}{\sigma\ \text{train}} \quad (1)$$

where μ_{train} and σ_{train} were computed exclusively from the training subset. Validation and test data were transformed using the same parameters to ensure consistency and avoid information leakage.

3.3 Problem Formulation
Let y_t denote monthly air temperature (T2M) at time t, and let u_t denote the vector of exogenous variables:

$$u_t = [RH2M_t,\ WS2M_t,\ RAD_t]. \quad (2)$$

Forecasting is formulated as a supervised learning problem. The objective is to estimate:

$$\hat{y}_{(t+h)} = f(\mathbf{H}_t), \quad (3)$$

where h∈{1,3} denotes the forecast horizon, and H_t contains all historical information available up to time t.
Two forecasting horizons were examined:
h = 1 month ahead
h = 3 months ahead
The multi-horizon setup allows evaluation of both short-term and medium-term predictive capacity.

3.4 Feature Engineering and Lag Structure
To capture temporal persistence and meteorological influences, structured lag vectors were constructed:

$$x_t = [y_{(t-1)}, \ldots, y_{(t-Ly)}, u_{(t-1)}, \ldots, u_{(t-Lu)}\ (4)]\ (4)$$

where:

$$Ly = 6. \quad Lu = 3. \quad (5)$$

In addition to lagged predictors, seasonal dynamics were modeled using Fourier seasonal components:

$$\sin\left(\frac{2\pi kt}{12}\right), \cos\left(\frac{2\pi kt}{12}\right), k = 1,2. \quad (6)$$

This method represents annual periodicity sparingly without the use of 11 dummy variables.
The correlation diagram between the temperature and the meteorological drivers is shown in Figure 4. Intermediate relationships warrant the exogenous predictors in the regression model.

Figure 4. Pearson correlation between monthly temperature and exogenous monthly meteorological variables.

3.5 Forecasting Models
3.5.1 Baseline Models
Two benchmark models were used:
Naïve Model

$$\hat{y}_{(t+h)} = y_t\ . \quad (7)$$

Seasonal Naïve Model

$$\hat{y}_{(t+h)} = y_{(t+h-12)}\,. \quad (8)$$

They are used as a lower bound of reference standards.

3.5.2 Regularised Regression Models

There were three regularized linear regression models in the test:

Ridge Regression

$$\mathbf{min}_{\beta}\ \Sigma_t\ (y_{(t+h)} - \beta_0 - \beta^T x_t)^2 + \lambda||\beta||_2^2\,. \quad (9)$$

Lasso Regression

$$\mathbf{min}_{\beta}\ \Sigma_t\ (y_{(t+h)} - \beta_0 - \beta^T x_t)^2 + \lambda||\beta||_1\,. \quad (10)$$

Elastic Net

$$\mathbf{min}_{\beta}\ \boldsymbol{\Sigma_t}\ (y_{(t+h)} - \beta_0 - \beta^T x_t)^2 + \lambda(\alpha||\beta||_1 + (1-\alpha)||\beta||_2^2). \quad (11)$$

Regularization parameters were chosen as a result of validation-set RMSE minimization.

Hyper parameter Selection

The regularisation parameters were selected using a systematic grid search strategy based on validation-set performance. For each model, a predefined grid of candidate values was explored:

- For Ridge and Lasso models, the regularisation parameter λ\lambdaλ was selected from a logarithmically spaced grid (e.g., λ∈ [10−4,102]).
- For Elastic Net, both λ and the mixing parameter α ∈ [0,1] were tuned simultaneously.

Model performance was evaluated using Root Mean Square Error (RMSE) on the validation set (2019–2021), and the optimal hyper parameters were chosen as those minimizing validation error.

To preserve temporal consistency and avoid information leakage, no random cross-validation was applied. Instead, model selection strictly followed the chronological validation framework.

3.6 Experimental Design

A strict chronological data splitting strategy was adopted to preserve temporal dependency and prevent information leakage:

- Training set: 2000–2018
- Validation set: 2019–2021
- Test set: 2022–2024

To ensure realistic out-of-sample evaluation, an expanding-window walk-forward forecasting framework was implemented during the testing phase. At each forecast origin time t:
The model was retrained using all available observations up to time t.
A forecast was generated for horizon h∈ {1,3}h, producing y_{t+h},
The training window was then expanded to include the next observation.
The process was repeated iteratively across the entire test period.

This approach ensures that:

- Only past information is used at each step.
- Model performance reflects true operational forecasting conditions.
- Evaluation is robust to temporal non-stationarity.

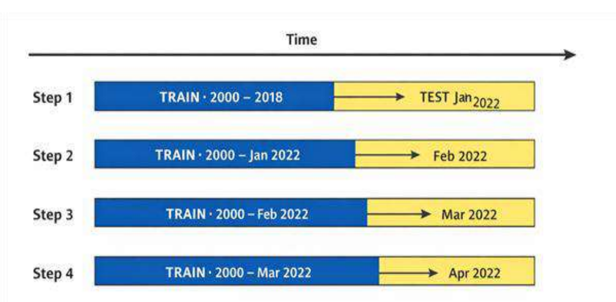

Fig.4: Expanding-window walk-forward validation scheme. The training set is progressively expanded over time, and forecasts are generated sequentially using only past observations at each step.

3.7 Evaluation Metrics

The accuracy of the forecast was determined using: Mean Absolute Error (MAE)

$$\mathbf{MAE} = (\frac{1}{N}) \Sigma |y_t - \hat{y}_t|. \quad (12)$$

Root Mean Square Error (RMSE)

$$\mathbf{RMSE} = \sqrt{\frac{1}{N}\sum_{t=1}^{N}(y_t - \hat{y}_t)^2} \,. \quad (13)$$

And can calculate Mean Absolute Percentage Error (MAPE) from:

$$\mathbf{MAP} = \frac{100}{N}\sum\nolimits_{t=1}^{N} |\frac{(y_t - \hat{y}_t)}{max(|y_t|, \varepsilon)}|. \quad (14)$$

3.8 Testing of Statistical Significance

The Diebold-Mariano (DM) test was used to compare the differentials in squared forecast errors:

$$\boldsymbol{d_t} = \boldsymbol{e^2_{(1,t)}} - \boldsymbol{e^2_{(2,t)}} \,, \quad (15)$$

and the difference in losses is given by:

$$\mathbf{DM} = \frac{\bar{d}}{\sqrt{\hat{\gamma}_d\,(0) \,/\, T}}. \quad (16)$$

The null hypothesis presupposes the same predictive accuracy. The level of significance was evaluated at 5 percent.
Figures 5 and 6 show the autocorrelation and partial autocorrelation structures and confirm the presence of persistence (up to a few monthly lags) and justify the lag specification of interest (Ly = 6, Lu = 3).

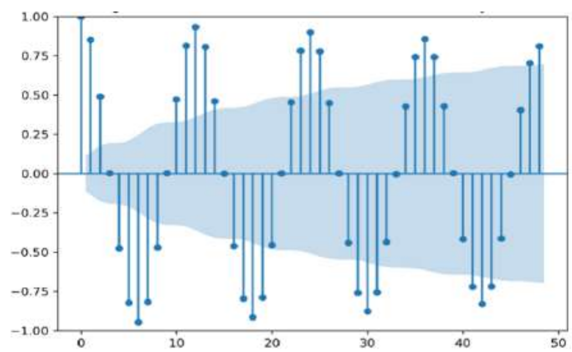

Fig.5: Monthly temperature Autocorrelation (ACF) of monthly temperature.

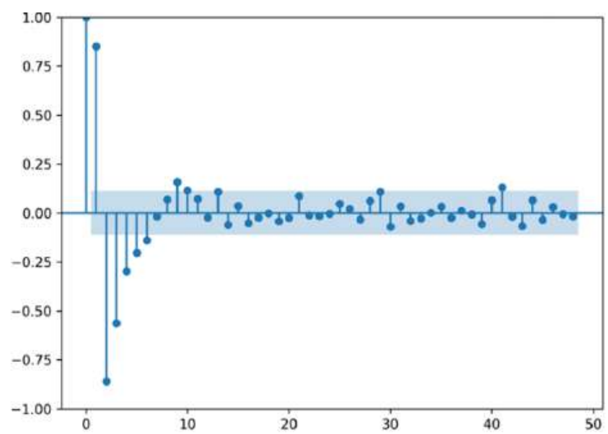

Fig.6: Partial correlation (PACF) of monthly temperature.

4. Results and Discussion

The results are presented in a structured manner, starting with short-term forecasting (h = 1), followed by medium-term forecasting (h = 3), to evaluate model performance across different prediction horizons systematically.

4.1 Forecasting Performance Evaluation

Forecasting performance was considered in a rigid chronological expanding-window walk-forward framework in the course of the independent test (2022-2024). Two forecasting horizons were considered, namely,

one-month ahead (h = 1) and

three-months ahead (h = 3).

One-month-ahead forecasting (h = 1).

Table 1 shows the comparative accuracy of forecasting by other competing models.

Table 1: Forecasting Performance (Test Period: 2022–2024, h = 1 Month)

Model	MAE (°C)	RMSE (°C)	MAPE (%)
Naïve (Persistence)	2.272	2.568	11.065
Seasonal Naïve	0.980	1.235	4.999
Ridge + Lags + Fourier	0.854	1.042	4.181
Lasso + Lags + Fourier	0.840	1.024	4.116
Elastic Net + Lags + Fourier	0.846	1.031	4.154

Fig. 8: Comparison between observed and predicted monthly temperature during the test period (2022–2024) for one-month-ahead forecasting. The regularised regression models closely follow the seasonal pattern with reduced short-term deviations.

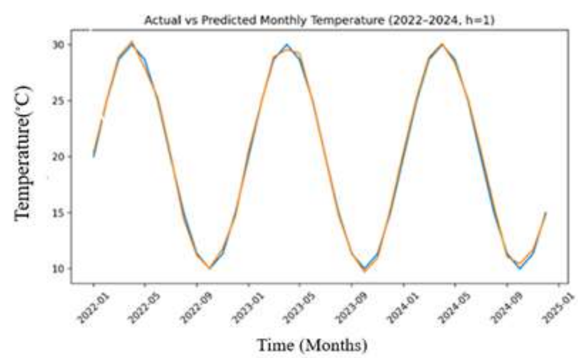

Fig. 7: Actual versus predicted monthly temperature (2022–2024, h = 1).

The regularised regression models are found to strongly follow the observed seasonal variations and understate the short-term variations.

The Lasso specification had the minimal RMSE (1.024 o C), which is a 60.1 percent reduction compared to the naive persistence model:

$$\frac{2.568 - 1.024}{2.568} \times 100 \approx 60.1\%$$

The improvement in comparison to the seasonal naive benchmark is:

$$\frac{1.235 - 1.024}{1.235} \times 100 \approx 17.1$$

Although the regularised regression models achieve lower RMSE values compared to the seasonal naïve benchmark, the Diebold–Mariano test indicates that these improvements are not statistically significant at the 5% level.

Three-Month-Ahead Forecasting (h = 3)

The three-month forecasting performance is obtained in Table 2.

Table2: Forecasting Performance (Test Period: 2022–2024, h = 3 Months)

Model	MAE (°C)	RMSE (°C)	MAPE (%)
Naïve (Persistence)	6.139	6.748	30.469
Seasonal Naïve	0.980	1.235	4.999
Ridge + Lags + Fourier	0.869	1.048	4.312
Lasso + Lags + Fourier	0.856	1.030	4.286
Elastic Net + Lags + Fourier	0.861	1.036	4.297

We can see that the persistence benchmark deteriorates substantially at longer horizons, and this is expected for strongly seasonal climate processes.

The Lasso model obtained an RMSE = 1.030°C, which expressed an 84.7 percent decrease compared to persistence:

$$\frac{6.748 - 1.030}{6.748} \times 100 \approx 84.7$$

When compared to the seasonal naïve model:

$$\frac{1.235-1.030}{1.235} \times 100 \approx 16.6$$

Notably, the performance is relatively constant (especially) in the range of h = 1 to h = 3 (1.024 to 1.030), indicating that the regression framework has a high level of structural robustness.

Fig. 8: Comparison of RMSE values across different forecasting models during the test period (2022–2024). Regularised regression models outperform the naïve benchmark, while their performance remains comparable to the seasonal naïve model, with no statistically significant differences observed.

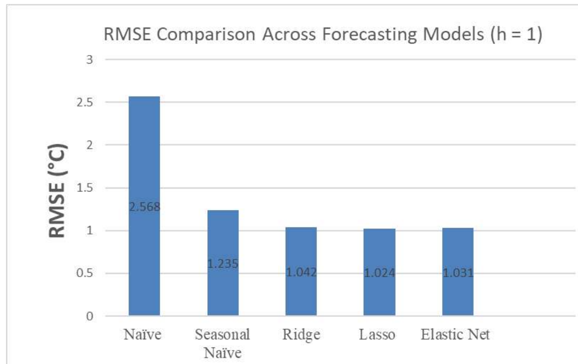

Fig. 8: Comparison of RMSE of different forecasting models (test time 2022 -2024).

4.2 Assessment of Statistical Significance.

The Diebold-Mariano (DM) test was used to test squared differentials in forecast errors to estimate whether the improvement is statistically significant.

h = 1 Month

Table 3: Diebold–Mariano Test Results (h = 1 Month)

Comparison	DM Statistic	p-value	Conclusion
Lasso vs Naïve	−5.666	2.13e−06	Significant
Lasso vs Seasonal Naïve	−1.406	0.169	Not Significant

h = 3 Months

Table 4: Diebold–Mariano Test Results (h = 3 Months)

Comparison	DM Statistic	p-value	Conclusion
Lasso vs Naïve	−4.683	4.17e−05	Significant
Lasso vs Seasonal Naïve	−0.798	0.430	Not Significant

The DM test approves statistically significant improvement of the advanced regression model over naïve persistence at both horizons ($p < 0.001$).

There are numerical improvements relative to the seasonal naïve model observed; these differences are not statistically significant at the 5% level. This outcome highlights the dominant strength of seasonal persistence in monthly temperature dynamics.

4.3 Interpretation of Findings

The lack of statistically significant improvement over the seasonal naïve benchmark can be attributed to the strong annual periodicity inherent in the temperature series. In highly seasonal climatic conditions, simple seasonal persistence captures a large portion of the variability, making it difficult for more complex models to achieve statistically significant gains.

This indicates that while the proposed regression framework provides consistent numerical improvements, the dominant seasonal structure limits the magnitude of statistically significant performance gains. Several major climatological and methodological lessons can be observed:

The annual periodicity of the Libyan temperature series is high enough, as shown by the competitive output of the seasonal naive benchmark.

It takes persistence, which is inadequate at longer horizons, leading to a large accumulation of errors.

Predictive accuracy is improved by the incorporation of exogenous atmospheric variables (humidity, wind speed, solar radiation).

The periodic dynamics of a Fourier seasonal structure are simple and not over-structured with too many parameters. With lasso regularization, generalization is improved because fewer of the weak or redundant lag coefficients are included in the model. There is evidence of stable performance over the different time horizons, indicating structural robustness instead of overfitting to short, fine-grained periods. The range of RMSE (1.02–1.05°C) calculated is within historical norms for monthly forecasted temperatures for semi-arid Mediterranean climates.

This highlights the importance of strong baseline models when evaluating forecasting performance in highly seasonal time series.

The applicability of the proposed framework is not limited to the study region. Given that many climatic regions exhibit similar seasonal patterns and dependence on exogenous atmospheric variables, the methodology can be extended to other geographical areas with comparable climatic characteristics.

However, the model performance may vary depending on the strength of seasonality, data quality, and the relevance of selected exogenous variables. In regions with weaker seasonal patterns or higher variability, the relative advantage of structured regression models over simple baseline approaches may become more pronounced.

Future work may explore model adaptation across diverse climatic zones and investigate the robustness of the proposed approach under different environmental conditions.

The findings of this study are consistent with previous applied research in climate time series forecasting. Several studies have shown that regularised regression models, particularly Lasso and Elastic Net, can achieve competitive predictive performance when combined with lagged features and exogenous meteorological variables. For instance, Zhang (2025) demonstrated that hybrid regression-based approaches can effectively capture temporal dependencies in environmental data, while Hajirahimi and Khashei (2019) highlighted the importance of combining structured models with external variables to improve forecasting accuracy.

In addition, applied forecasting studies, such as Athanasopoulos et al. (2011), have emphasized the importance of benchmark models, particularly in strongly seasonal datasets. Their findings indicate that simple models, including seasonal naïve approaches, can remain highly competitive due to the dominance of periodic patterns. This observation is consistent with the results of the present study, where improvements over the seasonal naïve benchmark were not statistically significant despite consistent numerical gains.

Furthermore, studies utilizing machine learning and statistical approaches for time series forecasting (e.g., Ahmed et al., 2022) have reported that while more complex models may provide marginal improvements, their advantage is often limited in highly structured seasonal environments. This reinforces the conclusion that model complexity does not always guarantee statistically significant performance gains.

4.4 Practical Implications

Dependable monthly weather prediction is important in:

- Energy demand estimation
- Agricultural scheduling
- Climate adaptation planning.
- Solar energy optimisation
- Urban heat risk management

The suggested regularised regression model provides:

- Interpretability
- Computational efficiency
- Operational deployability
- Constant of multi-horizon working stability.

Conclusion

This study developed a structured, regularized regression framework for monthly temperature forecasting in a semi-arid Mediterranean climate using long-term NASA POWER data.

The proposed approach, which integrates lag-based feature engineering, exogenous meteorological variables, and Fourier-based seasonal components, demonstrated consistent predictive performance across multiple forecasting horizons. The results indicate that regularised regression models, particularly Lasso, achieve substantial improvements over naïve persistence models.

However, while the proposed framework yields lower error metrics compared to the seasonal naïve benchmark, these improvements are not statistically significant according to the Diebold–Mariano test. This highlights the strong influence of seasonal periodicity in temperature dynamics, where simple seasonal persistence models can capture a large portion of variability.

Overall, the proposed methodology provides a robust, interpretable, and computationally efficient forecasting approach, offering competitive performance relative to baseline models without relying on complex nonlinear techniques.

Ethical Approval

This study is based exclusively on publicly available meteorological data obtained from the NASA POWER database. It does not involve human participants or animal subjects. Therefore, no ethical approval was required.

References

[1] M. A. Farahani et al., “Time-series forecasting in smart manufacturing systems: An experimental evaluation of state-of-the-art algorithms,” Reliable Computing in Industrial Manufacturing, 2025.

[2] X. Zhang, “Research on time series forecasting method based on ARIMA–ZKF prediction model,” Sustainability, vol. 17, no. 7, p. 2993, 2025.

[3] R. D. T. de Lima et al., “Time series forecasting with exogenous variables: A literature review and comparative perspectives,” Principia (IFPB), 2025.

[4] E. H. Houssein, “Artificial intelligence and classical statistical models for time series forecasting,” Journal of Big Data Analytics, 2025.

[5] H. Mahajan, “Comparative study of ARIMA, LSTM, and Prophet models,” Preprints, 2026.

[6] F. Yao et al., “CaReTS: A multi-task framework unifying classification and regression,” arXiv preprint.

[7] D. Shah and M. Thaker, “A review of time series forecasting methods,” Int. J. Research and Analytical Reviews, vol. 11, 2024.

[8] M. Kolambe and S. Arora, “Forecasting the future: A comprehensive review,” Journal of Electrical Systems, vol. 20, 2024.

[9] J. Kim et al., “A comprehensive survey of time series forecasting,” 2024.

[10] G. E. P. Box and G. M. Jenkins, “Time series analysis: Forecasting and control,” JRSS B, vol. 33, no. 2, 1971.

[11] C. Chatfield, Time-series forecasting. Chapman & Hall/CRC, 2000.

[12] D. C. Montgomery et al., Introduction to time series analysis and forecasting. Wiley, 2008.

[13] H. Akaike, “A new look at the statistical model identification,” IEEE Trans. Autom. Control, vol. 19, no. 6, 1974.

[14] P. Mondal et al., “Study of effectiveness of ARIMA in forecasting stock prices,” IJ CSEA, vol. 4, no. 2, 2014.

[15] Z. Hajirahimi and M. Khashei, “Hybrid structures in time series modeling,” Eng. Applications of AI, vol. 86, 2019.

[16] A. Tealab et al., “Forecasting of nonlinear time series using ANN,” Future Computing and Informatics Journal, vol. 2, 2017.

[17] P. H. B. de Melo et al., “COVID-19 forecasting using MLP,” Procedia Computer Science, vol. 181, 2021.

[18] V. N. Vapnik, The Nature of Statistical Learning Theory, 2nd ed. Springer, 2000.

[19] M. A. Ahamed and Q. Cheng, “TimeMachine: A time series is worth four Mambas,” arXiv preprint, 2024.

[20] D. M. Ahmed et al., "Deep sequential models for forecasting time series," Applied Computational Intelligence, 2022.

[21] J. M. L. Alcaraz and N. Strodthoff, "Diffusion-based time series forecasting," arXiv preprint, 2022.

[22] T. Alghamdi et al., "Forecasting traffic congestion using ARIMA," Proc. IWCMC, 2019.

[23] A. F. Ansari et al., "Chronos: Learning the language of time series," arXiv preprint, 2024.

[24] G. Athanasopoulos et al., "The tourism forecasting competition," Int. J. Forecasting, vol. 27, no. 3, 2011.

[25] A. Behrouz et al., "Chimera: Modeling multivariate time series," arXiv preprint, 2024.

[26] K. Benidis et al., "Deep learning for time series forecasting," ACM Computing Surveys, vol. 55, no. 6, 2022.