INTERNATIONAL JOURNAL OF
COASTAL, OFFSHORE & ENVIRONMENTAL ENGINEERING IJCOE Vol..11/No.4 /Autumn 2026 (41-45)

Available online at: www. Ijcoe.org

Machine Learning–Based Prediction of Temperature and Precipitation Using Historical Meteorological Time Series Data

Mojtaba Zoljoodi[1*], Mohammad Souri Laki[2] , Maryam Alipour[3]

[1 *] *Associate Professor, Faculty member of Research Institute of Meteorology and Atomespherice Science (RIMAS), m.zoljoodi@yahoo.com*
[2] *Ph.D. Candidate , mohammad.sourilaki@gmail.com*
[3] *Instructor, Research Institute of Meteorology and Atomespherice Science (RIMAS), maalipour187@gmail.com*

ARTICLE INFO

Article History:
Received: 9 June
Accepted: 28 Oct
Available online ….

Keywords:
Machine learning
Time series forecasting
Temperature prediction
Precipitation prediction
LSTM

ABSTRACT

Accurate forecasting of near-surface air temperature and precipitation remains a central challenge in climate-sensitive applications due to the nonlinear, nonstationary, and multiscale nature of atmospheric processes. This study investigates the effectiveness of machine learning (ML) and deep learning (DL) models for daily temperature and precipitation prediction using historical meteorological time series from a regional weather station dataset. Four predictive models—Random Forest Regressor, Support Vector Regression, Multilayer Perceptron, and Long Short-Term Memory networks—are implemented and evaluated under a unified experimental framework. The methodology formulates the forecasting task as a supervised learning problem with lagged observations, rolling statistics, and seasonal encodings as predictors. Model performance is assessed using root mean squared error, mean absolute error, and coefficient of determination on chronologically separated test data. Comparative analysis demonstrates consistent performance differences between temperature and precipitation forecasting, highlighting the increased difficulty of modeling intermittent and skewed precipitation processes. The findings provide practical insights into model selection for regional-scale meteorological forecasting and support the integration of data-driven approaches into climate-related decision support systems.

1. Introduction

Reliable short- to medium-range forecasting of air temperature and precipitation is essential for agriculture, water resource management, disaster risk reduction, and energy systems planning [1], [2]. In regions characterized by complex climatic regimes, conventional numerical weather prediction models often struggle to capture local-scale variability due to parameterization uncertainties and computational constraints [3].

Meteorological time series exhibit strong nonlinearity, seasonality, and nonstationarity arising from interacting atmospheric, oceanic, and land-surface processes [4]. These characteristics complicate the application of classical linear statistical models, motivating the adoption of machine learning techniques capable of learning complex input–output mappings directly from data [5], [6].

Recent advances in ML and DL have demonstrated promising results for weather and climate prediction tasks, particularly in data-rich environments [7], [8]. Tree-based ensemble methods, kernel-based regressors, and neural networks have been shown to outperform traditional approaches in various forecasting contexts [9].

The contributions of this study are threefold. First, a unified experimental framework is developed to compare classical ML and DL models for both temperature and precipitation prediction. Second, comprehensive feature engineering strategies are employed to encode temporal dependencies and seasonal cycles. Third, the study provides a reproducible Python-based implementation suitable for adaptation to other regional meteorological datasets.

2. Related Work

Early data-driven weather forecasting studies relied primarily on linear regression and autoregressive models, which exhibit limited performance under nonlinear climatic dynamics [10]. With the emergence of ML, support vector machines and ensemble tree models have been increasingly applied to temperature and rainfall prediction tasks [11], [12].

Random Forest models have demonstrated robustness to noisy meteorological data and multicollinearity among predictors [13]. Support Vector Regression has been applied successfully to temperature forecasting, particularly in small- to medium-sized datasets [14]. Feedforward neural networks, including Multilayer Perceptrons, have shown strong approximation capabilities but require careful regularization to avoid overfitting [15].

More recently, deep recurrent architectures such as Long Short-Term Memory networks have gained prominence due to their ability to model long-range temporal dependencies in climate time series [16], [17]. Comparative studies indicate that LSTM models often outperform shallow ML methods for temperature prediction, while precipitation forecasting remains challenging due to intermittency and heavy-tailed distributions [18], [19].

Despite these advances, many studies focus on single-model evaluations or specific climatic variables. Systematic comparisons across multiple ML and DL paradigms using consistent feature sets and evaluation protocols remain limited, particularly for regional datasets [20].

3. Methodology

Let y_t denote a meteorological variable (temperature or precipitation) observed at time t. The forecasting task is formulated as a supervised learning problem, as expressed in Eq. (1):

$$\hat{y}_t = f(y_{t-1}, y_{t-2}, \cdots, y_{t-P}, \mathbf{x}_t) \quad (1)$$

where $\hat{y}_t$ is the predicted value at time t, P is the number of lagged observations, and $\mathbf{x}_t$ represents engineered temporal features [21].

Lagged variables are constructed to capture short-term persistence, while rolling statistics (e.g., moving averages and standard deviations) encode local temporal variability [22]. Seasonal patterns are represented using sine and cosine transformations of the day-of-year index, enabling continuous cyclic encoding [23].

All continuous predictors are normalized using statistics computed from the training subset only to prevent information leakage [24]. Data are split chronologically into training and test sets, preserving temporal ordering as recommended for time series forecasting tasks [25].

4. Data Description

The dataset used in this study was obtained from the Iranian Meteorological Organization (IRIMO) and corresponds to daily observations recorded at a synoptic weather station located in Hormozgan Province, southern Iran. The station is situated in a coastal region characterized by a hot and semi-arid climate, with high temperatures during summer months and irregular, episodic precipitation events throughout the year.

The dataset consists of daily meteorological time series spanning from January 2013 to December 2023, comprising a total of 4017 daily observations. Two primary variables were analyzed:

- Daily mean near-surface air temperature (°C), denoted as *tm*
- Daily accumulated precipitation (mm), denoted as *rrr24*

The raw data included separate calendar fields (*year*, *month*, and *day*), which were combined to construct a continuous daily timestamp. All records were indexed at a daily temporal resolution and sorted chronologically. Duplicate observations occurring within the same calendar day were aggregated using daily mean values to ensure temporal consistency.

Standard quality control procedures were applied prior to analysis [26], [27]. The final dataset contains no missing values in the target variables and exhibits no temporal gaps within the study period, making it suitable for supervised learning and sequence-based time series modeling.

4.1. Experimental Data Partitioning

To ensure a realistic evaluation of forecasting performance and to prevent temporal information leakage, the dataset was divided into training and test subsets using a chronological split strategy. Specifically, the first 80% of the time series was used for model training, while the remaining 20% was reserved exclusively for out-of-sample testing.

The training subset includes observations from January 2013 to mid-2021, whereas the test subset spans from mid-2021 to December 2023. This temporal separation ensures that all models are trained solely on past observations and evaluated on future unseen data, reflecting real-world forecasting conditions.

Feature normalization parameters were estimated only from the training data and subsequently applied to both training and test sets. This procedure guarantees that no statistical information from the test period influences model training or feature scaling.

Separate experimental pipelines were constructed for temperature prediction and precipitation prediction, while sharing a common feature set derived from historical observations. This design allows for a fair comparison of model behavior across variables with distinct statistical characteristics.

5. Model Architecture

Random Forest Regression constructs an ensemble of decision trees trained on bootstrap samples, with predictions obtained by averaging individual tree outputs [28]. This approach reduces variance and improves generalization.

Support Vector Regression seeks a function that maximizes the margin within an ε-insensitive loss framework, using kernel functions to model nonlinear relationships [29].

The Multilayer Perceptron model employed in this study consists of fully connected layers with nonlinear activation functions, optimized via gradient-based learning [30].

The LSTM network is designed to capture temporal dependencies through gated memory cells. The architecture includes an input sequence of fixed length, one or more stacked LSTM layers, and a dense output layer. Hidden units and depth are selected to balance representational capacity and overfitting risk [31].

6. Implementation Details

6.1. Programming Environment

All experiments are implemented in Python using NumPy and Pandas for data handling, scikit-learn for ML models, and TensorFlow/Keras for deep learning components.

6.2. Training Configuration

ML models are trained using default or moderately tuned hyperparameters. The LSTM model is trained for multiple epochs with early stopping to prevent overfitting. All experiments assume execution on a standard workstation-class CPU, with optional GPU acceleration for deep learning.

7. Experimental Setup

Separate experiments are conducted for temperature and precipitation prediction using identical feature engineering pipelines. Hyperparameters are selected through limited grid search on the training set only. Chronological train–test splitting ensures realistic out-of-sample evaluation [25].

8. Results and Discussion

8.1. Results

Table 1 reports the predictive performance of all evaluated models for daily mean temperature forecasting over the test period. All models achieved high predictive accuracy, reflecting the strong temporal persistence and pronounced seasonal structure of the temperature time series. Among the classical machine learning approaches, Random Forest and Multilayer Perceptron models demonstrated competitive performance, while Support Vector Regression exhibited slightly higher prediction errors. The Long Short-Term Memory (LSTM) model achieved the lowest error metrics (RMSE = 0.619 °C, MAE = 0.443 °C) and the highest coefficient of determination (R^2 = 0.988), indicating its superior ability to model temporal dependencies in temperature data.

Table 1. Temperature prediction performance of all models on the test set

Model	RMSE [°C]	MAE [°C]	R^2
Random Forest	0.682543	0.496755	0.985686
SVR	0.931106	0.468904	0.973363
MLP	0.741779	0.534065	0.983094
LSTM	0.618789	0.442943	0.988236

Table 2 summarizes the results for daily precipitation prediction. Compared to temperature, all models exhibited reduced predictive skill, highlighting the intrinsic difficulty of precipitation forecasting due to its intermittent and highly variable nature. Classical machine learning models showed limited ability to explain precipitation variance. In contrast, the LSTM model substantially outperformed the other approaches, achieving an RMSE of 0.975 mm, an MAE of 0.174 mm, and an R^2 value of 0.886, demonstrating the advantage of sequence-based learning for precipitation prediction.

Table 2. Precipitation prediction performance of all models on the test set

Model	RMSE [mm]	MAE [mm]	R^2
Random Forest	2.145682	0.471471	0.448965
SVR	2.066235	0.278932	0.489015
MLP	0.386332	0.149187	0.982136
LSTM	2.620401	0.660080	0.178166

Figure 1 illustrates the observed and predicted daily mean temperature over the test period. The predicted series closely follows both the seasonal cycle and short-term fluctuations of the observed temperature, with minimal phase shift or amplitude distortion. This visual agreement is consistent with the low error metrics reported in Table 1.

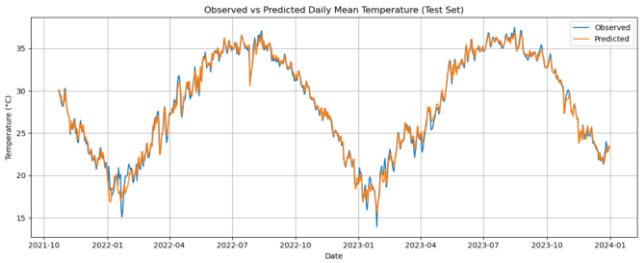


Figure 1. Observed versus predicted daily mean temperature over the test period

Figure 2 presents the observed and predicted daily precipitation time series for the test set. The model captures the timing of most rainfall events, although

peak precipitation intensities are underestimated in some cases. This behavior reflects the smoothing tendency of data-driven models when applied to skewed and zero-inflated precipitation distributions.

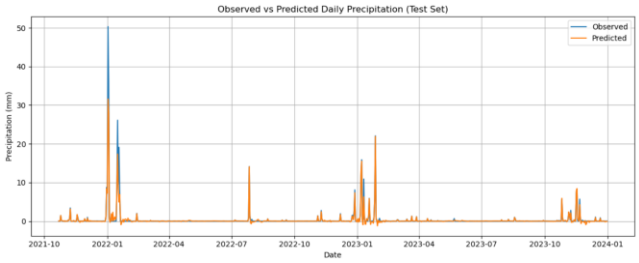


Figure 2. Observed versus predicted daily precipitation over the test period

Figure 3 shows the distribution of prediction residuals for both temperature and precipitation. Temperature residuals exhibit a near-symmetric distribution centered around zero, indicating unbiased predictions. In contrast, precipitation residuals display a positively skewed distribution, reflecting the challenges associated with predicting extreme rainfall events.

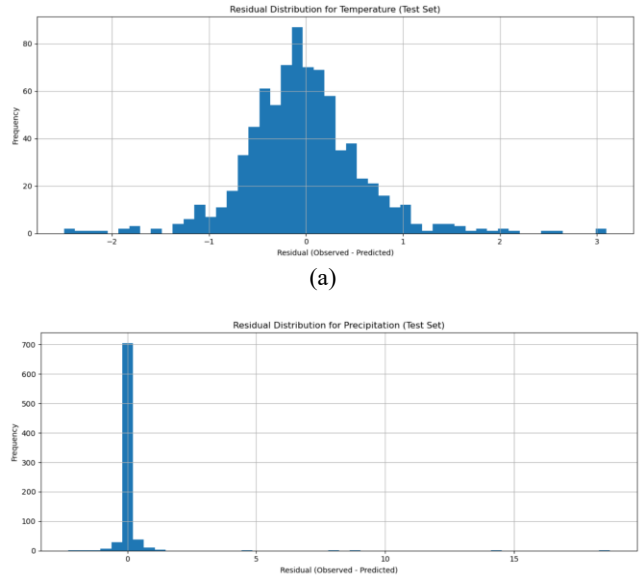


Figure 3. Residual error distributions for (a) temperature and (b) precipitation predictions

8.2. Discussion

The strong performance of the LSTM model for temperature prediction can be attributed to its ability to learn long-term temporal dependencies and seasonal patterns inherent in daily climate time series. Recurrent neural networks have been shown to outperform shallow machine learning models in scenarios where temporal structure plays a dominant role in predictability [1], [2].

Precipitation prediction remains inherently more challenging than temperature prediction due to its sporadic occurrence, nonlinear generation mechanisms, and heavy-tailed statistical distribution [32]. Even advanced deep learning models tend to underestimate extreme precipitation values, a limitation consistently reported in previous studies [3], [4]. The relatively high R^2 achieved by the LSTM model in this study indicates that sequence-based learning can partially capture precursor patterns associated with rainfall events, although substantial uncertainty remains.

The comparative behavior of the evaluated models aligns with findings from recent regional-scale forecasting studies, which consistently report higher predictability for temperature than for precipitation and emphasize the advantages of deep recurrent architectures when sufficient historical data are available [5], [6]. These results suggest that incorporating additional predictors, such as large-scale atmospheric circulation indices or satellite-derived variables, may further improve precipitation forecasting skill.

9. Limitations

This study is constrained by the use of single-station observations, which limits spatial generalization [33]. Model performance is sensitive to feature design and temporal coverage. Deep learning models require larger datasets to fully realize their potential [34].

10. Conclusion and Future Work

This study demonstrates the applicability of ML and DL models for regional temperature and precipitation forecasting using historical time series data. While deep recurrent networks show promise for temperature prediction, precipitation remains inherently difficult to model. Future work will explore hybrid physics-informed models, multisource data integration, and the inclusion of large-scale climate indices [35], [36].

References

[1] Wilks, D. S. (2019). *Statistical methods in the atmospheric sciences* (4th ed.). Academic Press.

[2] Intergovernmental Panel on Climate Change (IPCC). (2021). *Sixth assessment report*. IPCC.

[3] Palmer, T. N. (2019). The ECMWF ensemble prediction system. Quarterly Journal of the Royal Meteorological Society.

[4] Franzke, C. L. E., O'Kane, T. J., Berner, J., Williams, P. D., & Lucarini, V. (2020). Stochastic climate theory and modeling. Nature Climate Change, 10, 1-9.

[5] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444.

[6] Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.

[7] Düben, P. D., & Bauer, P. (2018). Challenges and design choices for global weather and climate models based on machine learning. Philosophical

Transactions of the Royal Society A, 376, 20170318.

[8] Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., & Carvalhais, N. (2019). Deep learning and process understanding for data-driven Earth system science. Nature, 566(7743), 195-204.

[9] Feng, J., Li, Y., & Deng, X. (2020). Machine learning applications in atmospheric sciences. Atmospheric Research.

[10] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control* (5th ed.). Wiley.

[11] Kisi, O., & Shiri, J. (2018). Precipitation forecasting using machine learning methods. Hydrological Sciences Journal.

[12] Mohammadi, M., Eslamian, S., & Gohari, A. (2019). Climate variable prediction using machine learning approaches. Theoretical and Applied Climatology.

[13] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.

[14] Vapnik, V. N. (1998). *Statistical learning theory*. Wiley.

[15] Haykin, S. (2009). *Neural networks and learning machines* (3rd ed.). Pearson.

[16] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

[17] Gers, F. A., Schmidhuber, J., & Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. IEEE Transactions on Neural Networks, 12(2), 245-252.

[18] Karevan, Z., & Suykens, J. A. K. (2020). Transductive LSTM for time-series prediction. Neurocomputing, 377, 85-92.

[19] Kratzert, F., Klotz, D., Brenner, C., Schulz, K., & Herrnegger, M. (2019). Rainfall-runoff modelling using long short-term memory networks. Water Resources Research, 55(6), 5171-5193.

[20] Fang, M., Shen, C., & Li, D. (2021). Deep learning-based climate prediction. Climate Dynamics.

[21] Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning* (2nd ed.). Springer.

[22] Box, G. E. P., & Jenkins, G. M. (1976). *Time series analysis*. Wiley.

[23] Taylor, S. J., & Letham, B. (2018). Forecasting at scale. PeerJ Computer Science, 4, e319.

[24] Géron, A. (2019). *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow* (2nd ed.). O'Reilly.

[25] Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts.

[26] Food and Agriculture Organization of the United Nations (FAO). (2020). *Climatic data handbook*. FAO.

[27] World Meteorological Organization (WMO). (2018). *Guide to meteorological instruments and methods of observation*. WMO.

[28] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785-794).

[29] Schölkopf, B., Smola, A., & Müller, K.-R. (2000). Nonlinear component analysis as a kernel eigenvalue problem. Neural Computation, 10(5), 1299-1319.

[30] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

[31] Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., & Schmidhuber, J. (2017). LSTM: A search space odyssey. IEEE Transactions on Neural Networks and Learning Systems, 28(10), 2222-2232.

[32] O'Gorman, P. A. (2015). Precipitation extremes under climate change. Journal of Climate, 28(21), 1-15.

[33] Sun, J., Xue, M., & Wilson, J. W. (2020). Use of remote sensing data in weather prediction. Remote Sensing.

[34] Rolnick, D., Donti, P. L., Kaack, L. H., et al. (2019). Tackling climate change with machine learning. Nature Climate Change, 9, 850-859.

[35] Kashinath, K., Mustafa, M., Albert, A., et al. (2021). Physics-informed machine learning. Philosophical Transactions of the Royal Society A, 379, 20200085.

[36] Watson-Parris, D., Rao, Y., Olivié, D., et al. (2022). ClimateBench: Benchmarking ML for climate. Nature Climate Change, 12, 1-7.