STAR Research Journal

Available online at www.starresearchjournal.com (Star International Journal)

RESEARCH

Seasonality-Aware CNN-GRU Modelling for Multi-Horizon Daily Weather Forecasting in South India

[1]PRAPILA S H, and [2]Dr. Sunil Damodar Rathod

[1]Research Scholar, Sunrise University, Alwar, Rajasthan
[2]Professor, Sunrise University, Alwar, Rajasthan

Abstract

Accurate short-range weather forecasting at regional scale remains difficult when weather variables are governed by both temporal continuity and seasonal transitions. This paper presents a seasonality-aware hybrid convolutional neural network-gated recurrent unit (CNN-GRU) model for multi-horizon daily weather forecasting in South India. The study used NASA POWER daily meteorological data from 2010 to 2024 for ten locations: Thiruvananthapuram, Kozhikode, Chennai, Coimbatore, Bengaluru, Mangaluru, Hyderabad, Visakhapatnam, Puducherry, and Madurai. The forecasting task involved four target variables: mean temperature, relative humidity, corrected precipitation, and wind speed. Forecasts were generated at 1-day, 3-day, and 7-day horizons using direct horizon-specific modelling. The proposed model combined one-dimensional convolutional feature extraction, GRU-based sequence learning, cyclical day-of-year encoding, and latitude-longitude location inputs. Independent testing was conducted on the unseen 2023-2024 period after validation-based look-back selection. A 30-day look-back window was selected for the 1-day horizon, whereas 60-day windows were selected for the 3-day and 7-day horizons. On the final test set, the seasonality-aware model achieved the lowest RMSE in eight of twelve variable-horizon comparisons and outperformed all benchmarks for every target variable at the 7-day horizon. At the 7-day horizon, it reduced RMSE relative to the conventional CNN-GRU ablation by 2.01% to 3.18% and relative to persistence by 16.79% to 26.97%. The results indicate that explicit seasonal representation is most useful when the forecast lead time increases and the predictive value of simple persistence weakens.

Keywords: *CNN-GRU, deep learning, multi-horizon forecasting, NASA POWER, seasonality, South India, time-series forecasting, weather prediction.*

I. INTRODUCTION

Weather forecasting supports decisions in agriculture, water management, energy planning, transportation, public safety, and disaster preparedness. In regions where weather is shaped by strong seasonal cycles, local circulation, coastal influence, and monsoon phases, forecasting models must account for both sequential dependence and annual periodicity. South India provides such a setting because its daily weather conditions differ across coastal, inland, humid, semi-arid, and elevated locations.

Conventional numerical weather prediction remains the operational foundation of modern forecasting, but data-driven methods have become important complementary tools. Deep learning models can learn nonlinear relationships from multivariate meteorological sequences and can be trained for location-specific or regional forecasting tasks. Recent advances in machine-learning weather prediction have shown that data-driven models can learn meaningful atmospheric patterns from historical data and reanalysis products [1], [2]. However, the value of simpler, transparent regional forecasting frameworks remains important, especially when the intended task is daily multi-variable forecasting for selected locations.

Meteorological time series are not only sequential but also seasonal. A model trained on daily weather values must interpret the same temperature, humidity, or wind condition differently depending on the time of year. A standard recurrent or convolutional sequence model may learn some seasonality implicitly, but explicit cyclical encoding can make the annual position of each observation available to the model. This is particularly relevant for South India because the seasonal structure includes winter, summer/pre-monsoon, southwest monsoon, and northeast/post-monsoon phases.

This paper develops and evaluates a seasonality-aware CNN-GRU model for 1-day, 3-day, and 7-day forecasting of mean temperature, relative humidity, corrected precipitation, and wind speed. The paper is written as a focused journal article rather than a thesis summary. Its main emphasis is the empirical value of explicit seasonal encoding within a hybrid CNN-GRU framework, evaluated against a conventional CNN-GRU ablation and a persistence baseline on an independent 2023-2024 test period.

The main contributions of the paper are as follows. First, it formulates a multi-location and multi-output daily weather-forecasting task for ten South Indian locations using a consistent open-data source. Second, it develops a hybrid CNN-GRU model that uses both cyclical seasonal features and static geographical inputs. Third, it evaluates forecast performance at multiple lead times using a chronological train-validation-test design. Fourth, it quantifies when seasonal encoding improves performance, showing that its advantage is most consistent at the 7-day horizon.

II. RELATED WORK AND RESEARCH POSITIONING

Deep learning has become a major modelling approach for sequential prediction because multilayer architectures can learn hierarchical representations from high-dimensional data [3]. In time-series forecasting, recurrent neural networks and their gated variants are frequently used to represent temporal dependencies. LSTM networks were introduced to address long-range dependency learning [4], while GRU-based recurrent units provide a simpler gated structure that is computationally efficient and effective for sequential representation [5].

Convolutional architectures are also useful for temporal modelling because one-dimensional filters can extract short local patterns from sequential inputs. Hybrid convolutional-recurrent structures are therefore suitable for meteorological data in which recent weather sequences contain both local fluctuations and longer temporal relationships. Related work on precipitation nowcasting and sequence modelling has shown the usefulness of convolutional recurrent ideas for weather-related prediction tasks [6].

Multi-step forecasting can be implemented through recursive, direct, or hybrid strategies. In a direct strategy, separate models are trained for specific forecast horizons, which avoids propagation of earlier prediction errors into later lead times [7]. This study adopted a direct multi-horizon strategy because it allowed 1-day, 3-day, and 7-day performance to be evaluated independently.

Model evaluation in forecasting must respect temporal order. Random splits are generally inappropriate for time-series forecasting because they can allow information from the future to influence model selection. Chronological splitting, final hold-out testing, and benchmark comparison provide stronger evidence of out-of-sample forecasting skill [8]. Forecast-error measures such as MAE and RMSE remain widely used, although interpretation should consider the variable, scale, and decision context [9].

Recent global systems such as GraphCast and GenCast have demonstrated the growing importance of machine learning for medium-range forecasting [1], [2]. However, regional daily forecasting remains a distinct research need, particularly when the aim is to develop reproducible models using open coordinate-based daily data. The present study addresses this need through a compact, regionally focused, seasonality-aware CNN-GRU model for South India.

III. MATERIALS AND METHODS

A. Study Area and Data Source

The study used daily meteorological data for ten selected locations in South India: Thiruvananthapuram, Kozhikode, Chennai, Coimbatore, Bengaluru, Mangaluru, Hyderabad, Visakhapatnam, Puducherry, and Madurai. The selected locations cover coastal and inland settings and include locations influenced by different seasonal rainfall regimes.

Daily data were obtained from the NASA POWER service, which provides solar and meteorological parameters for user-selected coordinates [10]. The study period covered 1

January 2010 to 31 December 2024. Each location contributed 5,479 daily records, giving 54,790 location-day records. The data were split chronologically: 2010-2021 for training, 2022 for validation, and 2023-2024 for independent testing.

TABLE I. DATASET STRUCTURE AND CHRONOLOGICAL PARTITIONING

Partition	Period	Purpose	Records	Locations
Training	2010-2021	Initial model fitting	43,830	10
Validation	2022	Look-back and configuration selection	3,650	10
Final testing	2023-2024	Independent out-of-sample evaluation	7,310	10
Complete dataset	2010-2024	Full data coverage	54,790	10

B. Input Features and Forecast Targets

The sequential meteorological inputs included maximum temperature, minimum temperature, mean temperature, relative humidity, corrected precipitation, wind speed, surface pressure, and solar radiation. The target outputs were mean temperature, relative humidity, corrected precipitation, and wind speed. Corrected precipitation was transformed as ln(1 + P) during model training and restored to millimetres per day before evaluation.

Seasonality was represented through cyclical day-of-year features. The day-of-year value was transformed into sine and cosine components so that the end and beginning of a calendar year were treated as adjacent positions in the annual cycle. Latitude and longitude were included as static location inputs. This design allowed the model to use both time-varying weather features and location-related information.

TABLE II. MODELLING VARIABLES

Category	Variables	Role
Sequential meteorological inputs	Maximum temperature, minimum temperature, mean temperature, relative humidity, corrected precipitation, wind speed, surface pressure, solar radiation	Historical predictors
Seasonal inputs	DOY_Sin and DOY_Cos	Cyclical annual representation
Static inputs	Latitude and longitude	Location representation
Forecast targets	Mean temperature, relative humidity, corrected precipitation, wind speed	Multi-output forecasts

C. Proposed Seasonality-Aware CNN-GRU Model

The proposed model used a hybrid deep-learning structure. A one-dimensional convolutional layer extracted local temporal patterns from each look-back window. The convolutional representation was passed to a GRU layer to learn sequential dependencies across the historical sequence. The static location branch processed latitude and longitude separately. The temporal and location representations were concatenated and passed through dense layers to generate four simultaneous forecasts.

A conventional CNN-GRU model was used as the main ablation benchmark. It retained the same general hybrid architecture but excluded the explicit seasonal features. This comparison allowed the contribution of DOY_Sin and DOY_Cos to be evaluated while keeping the core CNN-GRU structure comparable.

D. Forecasting Strategy and Evaluation

Separate direct models were developed for 1-day, 3-day, and 7-day horizons. Candidate look-back windows of 14, 30, and 60 days were evaluated using the 2022 validation data. The selected configurations were then refitted using the combined 2010-2022 development data before final testing on 2023-2024.

The model was compared with a persistence baseline and a conventional CNN-GRU ablation. MAE, RMSE, mean bias error (MBE), and coefficient of determination (R2) were used for evaluation. RMSE was treated as the principal criterion for identifying the best model in each variable-horizon comparison because it penalises larger errors more strongly than MAE [9].

IV. RESULTS

A. Validation-Based Look-Back Selection

The validation-stage experiment showed that the most suitable historical context depended on forecast horizon. A 30-day look-back window produced the lowest validation loss for the 1-day horizon, whereas 60-day windows produced the lowest validation loss for the 3-day and 7-day horizons.

TABLE III. VALIDATION-BASED LOOK-BACK SELECTION

Forecast horizon	Look-back tested	Validation loss	Best epoch	Selection
1 day(s)	30 days	0.118584	28	Selected
1 day(s)	14 days	0.123335	9	Not selected
1 day(s)	60 days	0.127992	5	Not selected
3 day(s)	60 days	0.216248	28	Selected

Forecast horizon	**Look-back tested**	**Validation loss**	**Best epoch**	**Selection**
3 day(s)	14 days	0.217446	22	Not selected
3 day(s)	30 days	0.221333	45	Not selected
7 day(s)	60 days	0.243961	23	Selected
7 day(s)	30 days	0.246604	15	Not selected
7 day(s)	14 days	0.250840	16	Not selected

B. Independent Final-Test Comparison

The final-test comparison was conducted on the unseen 2023-2024 data. The seasonality-aware CNN-GRU achieved the lowest RMSE in eight of the twelve variable-horizon combinations. The conventional CNN-GRU was best in three combinations, while persistence was best only for 1-day mean-temperature forecasting. This distribution of best results indicates that the proposed model provided the strongest overall test performance, while also showing that no model should be assumed superior for every weather variable and lead time.

The most consistent result appeared at the 7-day horizon. At this lead time, the proposed seasonality-aware model achieved the lowest RMSE for all four target variables. This suggests that explicit seasonal information became more valuable when the forecast target moved farther from the last observed day.

TABLE IV. FINAL-TEST RMSE COMPARISON ACROSS MODELS

Horizon	**Forecast variable**	**Seasonality-aware CNN-GRU**	**Conventional CNN-GRU**	**Persistence**	**Best model**
1 day(s)	Mean temperature	0.605	0.638	0.596	Persistence

Horizon	Forecast variable	Seasonality-aware CNN-GRU	Conventional CNN-GRU	Persistence	Best model
	(deg C)				
1 day(s)	Relative humidity (%)	3.819	3.770	3.788	Conventional
1 day(s)	Corrected precipitation (mm/day)	7.715	7.852	8.815	Seasonality-aware
1 day(s)	Wind speed (m/s)	0.532	0.536	0.575	Seasonality-aware
3 day(s)	Mean temperature (deg C)	0.887	0.874	0.944	Conventional
3 day(s)	Relative humidity (%)	5.369	5.407	5.952	Seasonality-aware
3 day(s)	Corrected precipitation (mm/day)	8.889	8.812	11.204	Conventional
3 day(s)	Wind speed (m/s)	0.761	0.776	0.934	Seasonality-aware
7 day(s)	Mean temperature (deg C)	1.016	1.049	1.231	Seasonality-aware
7 day(s)	Relative	6.063	6.188	7.287	Seasonality-

Horizon	**Forecast variable**	**Seasonality-aware CNN-GRU**	**Conventional CNN-GRU**	**Persistence**	**Best model**
	humidity (%)				aware
7 day(s)	Corrected precipitation (mm/day)	8.792	9.050	11.980	Seasonality-aware
7 day(s)	Wind speed (m/s)	0.818	0.842	1.120	Seasonality-aware

C. Effect of Seasonal Encoding at the 7-Day Horizon

The 7-day horizon was used to examine the practical effect of explicit seasonal encoding because the proposed model was best for all variables at this lead time. Relative to the conventional CNN-GRU ablation, RMSE reductions ranged from 2.01% for relative humidity to 3.18% for mean temperature. Relative to persistence, reductions ranged from 16.79% for relative humidity to 26.97% for wind speed.

TABLE V. RELATIVE RMSE REDUCTION OF THE PROPOSED MODEL AT THE 7-DAY HORIZON

Forecast variable	**Reduction vs. conventional CNN-GRU (%)**	**Reduction vs. persistence (%)**
Mean temperature (deg C)	3.18	17.48
Relative humidity (%)	2.01	16.79
Corrected precipitation (mm/day)	2.85	26.61
Wind speed (m/s)	2.93	26.97

D. Variable-Wise Interpretation

Mean temperature showed strong short-term continuity. Persistence produced the best 1-day RMSE for this variable, which indicates that the most recent observed value was highly informative for next-day temperature. At longer horizons, however, deep-learning models became more competitive, and the seasonality-aware model produced the best 7-day result.

Relative humidity benefited from the proposed model at the 3-day and 7-day horizons. Corrected precipitation remained the most difficult variable because rainfall was irregular, event-driven, and strongly affected by seasonal and local factors. Wind speed showed the most consistent advantage for the seasonality-aware model, which produced the lowest RMSE at all three horizons.

E. Location-Wise and Seasonal Patterns

Location-wise analysis showed that the proposed model did not perform uniformly across the ten cities. Average RMSE values differed by variable and location, indicating that regional forecasting skill was affected by local weather behaviour. For example, Kozhikode showed the lowest average RMSE for mean temperature and wind speed, Thiruvananthapuram for relative humidity, and Coimbatore for corrected precipitation. Hyderabad showed the highest average RMSE for mean temperature and relative humidity, Mangaluru for corrected precipitation, and Visakhapatnam for wind speed.

Seasonal analysis also showed variable-specific behaviour. The southwest monsoon period produced the highest average precipitation and wind-speed RMSE, while the summer/pre-monsoon period produced the highest average mean-temperature and relative-humidity RMSE. These findings indicate that a useful regional forecasting model should be evaluated by season rather than only through annual aggregate statistics.

TABLE VI. LOCATION-WISE AVERAGE RMSE OF THE PROPOSED MODEL ACROSS THREE HORIZONS

Location	Mean temp.	Relative humidity	Precipitation	Wind speed
Bengaluru	1.021	5.892	5.131	0.727
Chennai	0.807	4.446	10.982	0.778
Coimbatore	0.872	5.244	4.956	0.552
Hyderabad	1.286	7.931	5.669	0.727
Kozhikode	0.576	4.174	8.357	0.468
Madurai	0.964	5.056	5.586	0.554
Mangaluru	0.583	4.371	15.844	0.474
Puducherry	0.614	3.925	6.267	0.818
Thiruvananthapuram	0.581	3.741	8.789	0.760
Visakhapatnam	0.742	4.666	6.650	0.990

TABLE VII. SEASONAL-PERIOD AVERAGE RMSE OF THE PROPOSED MODEL ACROSS THREE HORIZONS

Season	Mean temp.	Relative humidity	Precipitation	Wind speed
Winter	0.835	5.279	7.595	0.669
Summer/pre-monsoon	1.022	6.509	4.454	0.629
Southwest monsoon	0.734	3.967	10.961	0.791
Northeast/post-	0.704	4.224	8.575	0.672

Season	**Mean temp.**	**Relative humidity**	**Precipitation**	**Wind speed**
monsoon				

V. DISCUSSION

The results provide three main insights. First, the usefulness of model complexity depends on forecast horizon. Persistence was difficult to beat for 1-day mean-temperature forecasting, but it weakened substantially at 3-day and 7-day horizons. This finding supports the inclusion of persistence as a mandatory baseline in regional weather-forecasting studies because it reveals whether a deep-learning model improves on the natural continuity of the series.

Second, explicit seasonality was most valuable at longer lead times. The seasonality-aware model did not dominate every 1-day or 3-day comparison, but it was consistently best at the 7-day horizon. This suggests that when immediate temporal continuity weakens, the annual position of the forecast target becomes a more useful source of information. The result is consistent with the seasonal structure of South Indian weather, where temperature, humidity, rainfall, and wind patterns are strongly conditioned by monsoon and pre-monsoon phases.

Third, the forecasting difficulty was variable-specific. Precipitation remained challenging even after logarithmic transformation and seasonal feature inclusion. Daily rainfall in the selected region has discontinuous occurrence, high spatial variability, and strong event dependence. This explains why precipitation produced larger errors and why further rainfall-specific modelling may be required. In contrast, wind speed showed the clearest model-based improvement, indicating that its daily behaviour contained learnable temporal and seasonal structure that the proposed architecture captured effectively.

The study also has a methodological implication. Regional weather models should not be evaluated only through a single pooled accuracy statistic. The location-wise and seasonal-period analyses showed that model performance can vary meaningfully across cities and climatic periods. For operational or advisory use, forecast users need information about which variables, locations, horizons, and seasons are associated with lower or higher expected error.

The paper has limitations. The dataset consisted of NASA POWER coordinate-based daily estimates rather than station-level observations. The modelling was restricted to daily forecasts and four target variables. The model did not include numerical weather-prediction outputs, synoptic indicators, sea-surface-temperature variables, or probabilistic uncertainty estimates. Therefore, the proposed model should be interpreted as a regional data-driven forecasting framework, not as a replacement for operational meteorological systems.

VI. CONCLUSION

This paper presented a seasonality-aware CNN-GRU framework for multi-horizon daily weather forecasting in South India. The model used multivariate daily meteorological sequences, cyclical day-of-year features, and static location inputs to forecast mean temperature, relative humidity, corrected precipitation, and wind speed at 1-day, 3-day, and 7-day horizons.

The independent final-test evaluation showed that the proposed model achieved the best RMSE in eight of twelve variable-horizon comparisons and was best for all four target variables at the 7-day horizon. The result demonstrates that explicit seasonal encoding is a useful addition to a hybrid CNN-GRU architecture, particularly when the forecast horizon increases. The strongest practical implication is that regional deep-learning weather models should include seasonality-aware feature representation and should be evaluated separately by horizon, variable, location, and season.

Future work can extend the framework by incorporating station observations, higher temporal resolution data, additional climatic predictors, numerical weather-prediction outputs, probabilistic forecasting, and advanced architectures such as attention-based, transformer-based, or ensemble models. Such extensions may be especially important for improving precipitation forecasting and for developing deployable regional decision-support systems.

Data Availability

The meteorological source data were obtained from the NASA POWER service. The processed datasets, model outputs, and prediction files used in this article were prepared as part of the associated research workflow.

REFERENCES

[1] R. Lam et al., "Learning skillful medium-range global weather forecasting," Science, vol. 382, no. 6677, pp. 1416-1421, 2023, doi: 10.1126/science.adi2336.

[2] I. Price et al., "Probabilistic weather forecasting with machine learning," Nature, vol. 637, pp. 84-90, 2025, doi: 10.1038/s41586-024-08252-9.

[3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436-444, 2015, doi: 10.1038/nature14539.

[4] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735-1780, 1997, doi: 10.1162/neco.1997.9.8.1735.

[5] K. Cho, B. van Merrienboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in Proc. 2014 Conf. Empirical Methods in Natural Language Processing (EMNLP), Doha, Qatar, 2014, pp. 1724-1734, doi: 10.3115/v1/D14-1179.

[6] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-C. Woo, "Convolutional LSTM network: A machine learning approach for precipitation nowcasting," in Advances in Neural Information Processing Systems 28, 2015, pp. 802-810.

[7] S. B. Taieb, G. Bontempi, A. F. Atiya, and A. Sorjamaa, "A review and comparison of strategies for multi-step ahead time series

forecasting based on the NN5 forecasting competition," Expert Systems with Applications, vol. 39, no. 8, pp. 7067-7083, 2012, doi: 10.1016/j.eswa.2012.01.039.

[8] V. Cerqueira, L. Torgo, and I. Mozetic, "Evaluating time series forecasting models: An empirical study on performance estimation methods," Machine Learning, vol. 109, pp. 1997-2028, 2020, doi: 10.1007/s10994-020-05910-7.

[9] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," International Journal of Forecasting, vol. 22, no. 4, pp. 679-688, 2006, doi: 10.1016/j.ijforecast.2006.03.001.

[10] National Aeronautics and Space Administration, "Prediction Of Worldwide Energy Resources (POWER) project," NASA Langley Research Center. [Online]. Available: https://power.larc.nasa.gov/. Accessed: Jul. 8, 2026.

[11] B. Lim and S. Zohren, "Time-series forecasting with deep learning: A survey," Philosophical Transactions of the Royal Society A, vol. 379, no. 2194, Art. no. 20200209, 2021, doi: 10.1098/rsta.2020.0209.

[12] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," Proceedings of the IEEE, vol. 86, no. 11, pp. 2278-2324, 1998, doi: 10.1109/5.726791.

[13] C. O. de Burgh-Day and T. Leeuwenburg, "Machine learning for numerical weather and climate modelling: A review," Geoscientific Model Development, vol. 16, pp. 6433-6477, 2023, doi: 10.5194/gmd-16-6433-2023.

[14] R. J. Hyndman and G. Athanasopoulos, Forecasting: Principles and Practice, 3rd ed. Melbourne, Australia: OTexts, 2021.