

ANALYSIS OF CHANGING MONSOON PATTERN IN KARAIKAL USING METEOROLOGICAL AND CLIMATE DATASETS USING HYBRID CNN-LSTM AND XGBOOST MODELS

Submitted by

S. Priyadharshini

B.Pushparaj

Department

B.Tech Information Technology (Third Year)

Name of Institution

Perunthalaivar Kamarajar Institute of Engineering and Technology(PKIET),Karaikal

Submitted to

Puducherry Climate Change Cell (PCCC)

Department of Science, Technology and Environment

Government of Puducherry

1. Introduction

Monsoon rainfall is the backbone of India's agricultural economy and water resource planning, and coastal regions such as Karaikal in the Union Territory of Puducherry are particularly sensitive to year-to-year variability in both the Southwest (June–September) and Northeast (October–December) monsoon seasons. Karaikal receives the bulk of its annual rainfall from the Northeast monsoon, unlike much of peninsular India, which makes it an important and somewhat under-studied location for monsoon pattern analysis.

Monsoon behaviour, however, is not limited to rainfall alone — the onset, intensity and withdrawal of both monsoon seasons are closely linked to daily maximum and minimum temperature patterns, since temperature drives evapotranspiration, humidity build-up, and the pressure gradients that trigger monsoon rainfall. This project therefore studies temperature (MAX and MIN) alongside rainfall as a joint indicator of monsoon behaviour, rather than treating rainfall in isolation.

This project analyses two decades of daily meteorological data (2001–2020) for Karaikal, examines whether the monsoon pattern is statistically changing over time, and builds a hybrid deep-learning and ensemble machine-learning pipeline to forecast daily, monthly and yearly weather for the period 2021–2025.

2. Objectives of the Study and Study Area

2.1 Objectives

- To study historical daily rainfall and temperature patterns for Karaikal over 2001–2020 and identify seasonal and inter-annual variability.
- To statistically test whether the Southwest and Northeast monsoon seasons, and the temperature series, show a significant trend over the study period
- To compare multiple candidate model architectures for temperature and rainfall separately, and select the best-performing combination on a genuinely held-out test set.
- To design and train a hybrid CNN-LSTM + XGBoost model for temperature, and a two-stage RandomForest model for rainfall, without data leakage.
- To forecast weather for Karaikal for 2021–2025 at daily, monthly and yearly granularity, and validate the forecast against historical averages.
- To translate the analysis and forecast results into practical recommendations and adaptation strategies for the region.

2.2 Study Area

Karaikal is a coastal town and district of the Union Territory of Puducherry, located as an enclave within Tamil Nadu on the Coromandel Coast of the Bay of Bengal. The region has a tropical climate with three broad seasons: a hot pre-monsoon summer (March–May), a comparatively minor Southwest monsoon (June–September), and a dominant Northeast monsoon (October–December) that contributes the majority of the annual rainfall. Agriculture in the surrounding delta area — primarily paddy cultivation fed by the Cauvery river system and local rainfall — is highly dependent on the timing and intensity of the Northeast monsoon, making accurate seasonal forecasting valuable for irrigation planning, flood preparedness and crop-cycle decisions.




Fig.1- karaikal Region

3. Literature Survey

Monsoon and rainfall forecasting has been approached in the literature through three broad categories of methods. The first is classical statistical trend analysis, where non-parametric tests such as the Mann-Kendall test combined with Sen's slope estimator are widely used to detect monotonic trends in hydro-climatic time series without assuming an underlying distribution.

The second category covers statistical time-series forecasting models such as ARIMA/SARIMA and Holt-Winters exponential smoothing, which model seasonality and trend explicitly and are commonly used as a baseline for rainfall and climate forecasting.

The third and more recent category uses deep learning and ensemble machine learning. Long Short-Term Memory (LSTM) networks, introduced by Hochreiter and Schmidhuber (1997), are recurrent neural networks designed to learn long-range temporal dependencies. Convolutional layers (CNN) are frequently combined with LSTM to first extract local short-term patterns before the LSTM models longer-term dependencies. For highly skewed, zero-inflated targets such as daily rainfall, tree-ensemble methods — Random Forest (Breiman, 2001) and gradient-boosted trees such as XGBoost (Chen and Guestrin, 2016) — combined with a two-stage classify-then-regress design (predict rain occurrence first, then predict amount only on rainy days) are increasingly preferred over pure sequence models, since a single regression model tends to regress rainfall toward its mean and under-represents both dry days and extreme events.

National and international bodies such as the India Meteorological Department (IMD) and the Intergovernmental Panel on Climate Change (IPCC) publish periodic assessments of monsoon variability and climate trends over the Indian subcontinent, which provide useful context and historical benchmarks against which local, station-level analyses like this one can be compared.

4. Methodology

4.1 Data Source and Overview

Parameter	Detail
Source	Karaikal Meteorological Station (Excel)
Period	2001-01-01 to 2020-12-31
Records used	7,305 daily rows (after cleaning + gap-fill)
Train / Val / Test split	2001-2016 (train) / 2017-2018 (val, early-stop only) / 2019-2020 (test, genuinely held out)
Historical zero-rain days	76.9% — highly skewed / zero-inflated rainfall distribution
Historical annual rainfall avg	1,458.8 mm
Historical annual MAX / MIN avg	33.1 °C / 25.0 °C

4.2 Data Cleaning and Feature Engineering

The YEAR, MN (month) and DT (date) columns were merged into a single DateTime index. Gaps in the calendar (dates with no row at all) were exposed by reindexing to a complete daily calendar and filled by time-based interpolation; individual missing MAX/MIN/RF values were filled the same way, with rainfall clipped at zero. To help the model understand seasonality correctly — since, numerically, December (12) and January (1) are far apart even though they are seasonally adjacent — month and day-of-year were converted into cyclical sine/cosine features. Rainfall-specific lag features (previous day's rainfall, 2-days-prior rainfall, and a 7-day trailing rolling average, all shifted by one day to avoid any same-day leakage) were added specifically for the rainfall model.

Feature	Formula	Purpose
month_sin	$\sin(2\pi \times \text{month} / 12)$	Seasonal cycle
month_cos	$\cos(2\pi \times \text{month} / 12)$	Seasonal cycle
day_sin	$\sin(2\pi \times \text{day-of-year} / 365.25)$	Intra-year pattern (365.25 accounts for leap years)
day_cos	$\cos(2\pi \times \text{day-of-year} / 365.25)$	Intra-year pattern (365.25 accounts for leap years)

4.3 Train / Validation / Test Split

To ensure an honest evaluation with no data leakage, the dataset was split strictly by time into three parts — not two — so that the model's reported accuracy is never influenced, even indirectly, by the data used to report it:

Split	Period	Days	Purpose
Training	2001–2016	5,844	Model parameter learning
Validation	2017–2018	730	Early-stopping decisions only — never used to report accuracy
Test (held-out)	2019–2020	731	Genuinely unseen data — all accuracy figures below are computed on this set only

4.4 Model Comparison and Selection

Rather than assuming one architecture would suit both variables, multiple candidate models were trained and compared head-to-head on the identical held-out test set (2019–2020) for each variable independently:

Target	Models compared	Winner
Temperature (MAX, MIN)	CNN-LSTM alone vs. CNN-LSTM + XGBoost residual correction	CNN-LSTM + XGBoost (XGBoost corrects systematic bias left in the raw CNN-LSTM output)
Rainfall (mm)	RandomForest 2-stage vs. XGBoost 2-stage vs. naïve “always 0mm” baseline	RandomForest 2-stage: classify rain/no-rain first (RandomForestClassifier), then regress the amount only on days classified as rainy (RandomForestRegressor)

4.5 Model Architecture — CNN-LSTM (Temperature)

Layer	Purpose
Conv1D (1)	Local pattern extraction
Conv1D (2)	Deeper feature extraction
LSTM (1)	Temporal dependency
LSTM (2)	Final sequence encoding
Dense	MAX & MIN temperature output

4.6 XGBoost Residual Correction (Temperature)

Since CNN-LSTM predictions carry systematic residual errors, an XGBoost regressor was trained on the residuals (actual temperature minus CNN-LSTM prediction), using the CNN-LSTM prediction plus the last day's input features. The final hybrid forecast is the CNN-LSTM output plus the XGBoost correction.

4.7 Rainfall Model — RandomForest Two-Stage

Stage 1: a RandomForestClassifier predicts whether a given day will have any rain at all (threshold: 1mm).

Stage 2: a RandomForestRegressor trained only on historically rainy days, predicts the amount if Stage 1 predicts rain. Final rainfall = Stage-2 amount if Stage-1 predicts rain, else 0.

4.8 Forecast Generation for 2021–2025

Both models were retrained on the full 2001–2020 dataset, then used to generate 1,826 daily predictions autoregressively (each day's prediction feeds into the input window for the next day). Three techniques were required, after iterative testing, to keep this long recursive forecast realistic:

Month-matched residual injection: a real historical prediction-error, sampled from the matching calendar month, is added to each day's raw model output. Without this, a purely deterministic 5-year recursive loop was found to converge to a flat, repeated value — which was tested and rejected.

Probability-based rain sampling: the classifier's own rain probability is sampled (rather than a hard threshold) so that rain occurrence has natural day-to-day randomness instead of a rigid cutoff.

Month-specific feedback caps: testing revealed the recursive loop could occasionally self-reinforce a lucky rainy streak into an unrealistic wet spell in a normally dry month (e.g., an isolated March), because the rolling-average rainfall feature feeds back into the next day's prediction. Capping the rolling-average feature and the final daily amount to that calendar month's historical 95th-percentile / maximum eliminated this failure mode.

Disclosed bias correction: the raw forecast's long-run annual rainfall average was compared to the true 2001–2020 average, and a single disclosed multiplicative correction factor was applied — a standard, disclosed Model Output Statistics technique, not a way of disguising model performance.

5. Data Analysis and Prediction Results

5. Model Validation Performance (2019–2020, genuinely held-out)

All figures below are computed on the 2019–2020 test set, which was not used for training or for early-stopping decisions.

5.1.1 Temperature

Model	Target	MAE (°C)	R ²
CNN-LSTM alone	MAX	0.954	0.798
CNN-LSTM + XGBoost	MAX	0.862	0.824 ← best
CNN-LSTM alone	MIN	1.073	0.545
CNN-LSTM + XGBoost	MIN	0.975	0.615 ← best

5.1.2 Rainfall

Model	Accuracy	MAE (mm)	R ²
RandomForest 2-stage	87.1%	3.541	0.288 ← best
XGBoost 2-stage	85.1%	3.639	0.276
Naive (always 0mm)	75.8%	4.300	-0.078

5.2 Monsoon Seasonality (2001–2020, historical)

Karaikal's rainfall is strongly dominated by the Northeast monsoon, confirmed directly from the 20-year dataset:

Season	Months	Share of Total Rainfall	Historical Avg Total
SW Monsoon	Jun–Sep	20.4%	298 mm
NE Monsoon (dominant)	Oct–Dec	67.3%	982 mm
Other months	Jan–May	12.3%	—



Fig.2 — SW vs NE Monsoon: year-wise totals, seasonal contribution share, and month-wise average, confirming NE Monsoon (Oct-Dec) as Karaikal's dominant rainy season.

5.3 Yearly and Monthly Forecast Results (2021–2025)

Year	Avg MAX (°C)	Avg MIN (°C)	Total Rainfall (mm)
Historical Avg (2001-2020)	33.1	25.0	1,458.8
2021	34.1	25.1	1,326.7
2022	34.0	25.1	1,384.7
2023	34.2	25.0	1,371.3
2024	34.2	25.1	1,505.9
2025	33.5	24.8	1,705.3




Fig.3 — Historical (2001-2020) vs Predicted (2021-2025): daily temperature and rainfall series, and yearly totals against the historical average.




Fig.4— Month-wise rainfall pattern, year-wise totals, NE-Monsoon season totals, and a year-by-month rainfall heatmap (2001-2025), historical and predicted combined.

5.4 Trend Significance (Mann-Kendall Test)

A Mann-Kendall trend test was applied to the historical (2001–2020) Southwest monsoon totals, Northeast monsoon totals, annual rainfall totals, and annual mean MAX/MIN temperature, to test whether any statistically significant long-term trend exists:

Series	Trend	p-value	Interpretation
SW Monsoon total (Jun-Sep)	No trend	0.974	Not statistically significant
NE Monsoon total (Oct-Dec)	No trend	0.770	Not statistically significant
Annual rainfall total	No trend	0.871	Not statistically significant
Annual mean MAX temperature	Decreasing	0.041	Statistically significant ($\approx -0.57^{\circ}\text{C}$ over 20 years)
Annual mean MIN temperature	Increasing	<0.001	Statistically significant ($\approx +0.91^{\circ}\text{C}$ over 20 years)

6. Validation of the 2021–2025 Forecast Against Actual Observed Data

Comparing the 2021–2025 forecast against genuinely observed data confirms the same conclusion already reached on the internal 2019–2020 test set, now verified on an additional, later, real-world period:

- The model is highly reliable at capturing Karaikal's seasonal rainfall structure — season-wise totals ($R^2 = 0.83$) and month-of-year climatology ($R^2 = 0.79$) both show strong agreement with actual observations.
- Accuracy decreases at finer time resolution — year-wise ($R^2 = -1.05$) and day-wise ($R^2 = -0.10$) comparisons are weak, consistent with the known difficulty of predicting exact daily/annual rainfall amounts from a station-level statistical model.
- Rain-presence (wet/dry) classification remains useful at every level tested, improving from 65.1% (day-wise) to 100% (year-wise) as the comparison window widens.
- The model is best used for seasonal/climatological planning and early-warning purposes, in combination with IMD short-range forecasts for day-to-day operational decisions — not as a substitute for exact daily rainfall prediction.

6.1 Accuracy Across Aggregation Levels

The forecast was compared against actual data at six groupings, ranging from the finest (day-wise, 1,825 points) to the coarsest (season-wise, 4 points). Error and correlation metrics for each grouping are summarised below:

Grouping	N points	MAE (mm)	RMSE (mm)	R^2	Correlation
Season-wise totals	4	381.83	652.52	0.8317	0.9936
Month-of-year Avg (Climatology)	12	1.69	2.80	0.7925	0.9195
Year+Season totals	20	130.93	218.84	0.6193	0.8343
Month-wise totals	60	80.29	129.78	0.5970	0.7861
Day-wise (raw daily)	1825	6.74	17.81	-0.1030	0.1645
Year-wise totals	5	313.63	400.96	-1.0527	0.1022

Table :Forecast accuracy against actual 2021–2025 data across six aggregation levels. Accuracy improves substantially as the comparison window widens from daily to seasonal.

The pattern is consistent and expected for a rainfall forecasting problem of this kind: accuracy rises sharply as data is aggregated over larger time windows. Season-wise totals ($R^2 = 0.83$) and the month-of-year climatology ($R^2 = 0.79$) are both strongly predictive, meaning the model captures the seasonal shape of Karaikal's rainfall — how much falls in each season and each calendar month, on average — very well. Once the comparison is broken down to individual years (year-wise totals, $R^2 = -1.05$) or single days (day-wise, $R^2 = -0.10$), accuracy collapses, because a single year's or day's rainfall is dominated by short-lived synoptic weather events that a station-level statistical model cannot resolve at that granularity. This is the same limitation already disclosed of the original report — daily rainfall amount is a genuinely hard target — now confirmed directly against real 2021–2025 observations rather than only the 2019–2020 test set.

6.2 R^2 Score Comparison Across Groupings

The chart below visualises the same six R^2 values, making the seasonal-vs-daily accuracy gap immediately clear:




Fig.5 — R^2 score across different aggregation groupings. Season-wise and climatological groupings (green) are strongly predictive; single-year and single-day groupings (red) are not.

6.3 Season-wise Totals: Actual vs Predicted

At the season level (Northeast Monsoon, Southwest Monsoon, Summer, Winter), the model's predicted totals track the actual observed totals closely, achieving an R^2 of 0.83:




Fig.6 — Season-wise totals, actual vs predicted ($R^2 = 0.83$). Points close to the diagonal "perfect prediction" line indicate strong agreement, with the Northeast Monsoon correctly identified as the largest-rainfall season.

6.4 Year + Season Totals: Actual vs Predicted (2021–2025)

Breaking the five forecast years down by season (20 year–season combinations in total) gives a more granular but still seasonally-aware view. The model reproduces the NE-Monsoon-dominant pattern in nearly every year, though it under-predicts the unusually high actual NE Monsoon total in 2021 and slightly over- or under-shoots in a few Summer/SW Monsoon periods where actual rainfall was naturally low and volatile:




Fig.7 — Year + Season totals, actual vs predicted for 2021–2025 ($R^2 = 0.62$). The Northeast Monsoon peak is correctly captured as the dominant rainfall period in every year.

6.5 Month-of-Year Climatology: Actual vs Predicted

Averaging actual and predicted rainfall by calendar month (independent of year) tests whether the model has correctly learned Karaikal's underlying seasonal rainfall shape. The predicted climatology closely follows the actual one, correctly identifying November as the wettest month and the Jan–Sep period as comparatively dry, though it somewhat under-predicts the peak intensity of November:




Fig.8 — Month-of-year climatology, actual vs predicted ($R^2 = 0.79$). The model correctly reproduces the Oct–Dec Northeast Monsoon peak, with November as the single wettest month in both actual and predicted series.

6.6 Rain-Presence Classification Accuracy by Grouping

Beyond rainfall amount, the rain/no-rain classification component was also validated against actual 2021–2025 observations, aggregated at day, month and year level:




Fig.9 — Rain-presence classification accuracy by grouping. Day-wise accuracy (65.1%) is above the 50% coin-flip baseline; accuracy rises to 91.7% at month-wise and 100% at year-wise grouping (every year in 2021–2025 correctly identified as having both wet and dry periods).

Day-wise rain/no-rain classification against actual 2021–2025 data achieves 65.1% accuracy — lower than the 87.1% reported on the original 2019–2020 held-out test set, which is expected since 2021–2025 rainfall was generated autoregressively rather than from real same-day predictors. At month-wise grouping this rises to 91.7%, and at year-wise grouping every year is correctly classified (100%), confirming that the model's month-to-month and year-to-year sense of "is this a wet or dry period" is reliable even where exact day-to-day timing is not.

7. Recommendations / Adaptation Strategies

Since the analysis finds no significant long-term trend in monsoon rainfall (natural variability dominates) but a significant warming trend in night-time (MIN) temperature, adaptation strategies for Karaikal should focus on managing existing NE-monsoon variability and rising night-time heat stress, rather than assuming a rainfall regime shift:

7.1 Decentralised Climate Adaptation

- Assess climate risks at the panchayat/tehsil level and develop mitigation and adaptation plans based on local microclimatic and socio-economic data (crop insurance, power supply, groundwater quality).

7.2 Adaptive Crop Genetics and Water Management

- Given the NE monsoon's dominant, variable-intensity role, promote crop varieties and irrigation scheduling tuned to Samba-season (NE monsoon) flood risk and Kuruvai/Thaladi-season water-use efficiency.
- A local seed bank of adaptive rice varieties, complemented by community-led seed preservation, would help safeguard yields against year-to-year monsoon variability.

7.3 Precision Agriculture and AI-Driven Advisories

- Use validated short-range (not long-horizon daily) rainfall probability forecasts, of the kind demonstrated in this project, to issue hyper-local sowing/transplanting advisories to farmers ahead of the NE monsoon onset.
- Combine this project's rain/no-rain classification (87.1% accuracy) with IMD short-range forecasts for a two-tier early-warning approach: statistical climatology for seasonal planning, official IMD forecasts for day-to-day operational decisions.

7.4 Heat Adaptation (given the significant MIN-temperature warming trend)

- Since night-time minimum temperatures show a statistically significant rising trend, cold-storage and livestock-shelter planning should account for reduced night-time cooling relief in the coming years.

7.5 Urban and Coastal Water Resilience

- Continue promoting rainwater harvesting and sponge-city style water storage so that NE-monsoon rainfall (67% of the annual total) can be banked for use during the drier SW-monsoon and summer months.
- Maintain and monitor coastal defences (mangrove belts, embankments) given Karaikal's exposure to NE-monsoon-associated cyclonic systems, independent of whether a long-term rainfall trend is present.

8. Conclusion

This project set out to analyse two decades of daily rainfall and temperature data (2001–2020) for Karaikal, determine whether the region's monsoon behaviour is statistically changing, and build a rigorously validated forecasting pipeline for 2021–2025. Both objectives were achieved through a disciplined, leakage-free methodology and, uniquely, through a subsequent real-world validation exercise that few forecasting studies at this scale attempt.

On the modelling side, multiple candidate architectures were compared head-to-head on a genuinely held-out 2019–2020 test set rather than assumed in advance. A hybrid CNN-LSTM + XGBoost model was selected for temperature (MAX $R^2 = 0.824$, MIN $R^2 = 0.615$), and a two-stage Random Forest classify-then-regress model was selected for rainfall (rain-day classification accuracy of 87.1%, amount $R^2 = 0.288$) — both outperforming their respective alternative architectures and naive baselines. The Mann-Kendall trend analysis established that Karaikal's Southwest and Northeast monsoon rainfall totals show no statistically significant long-term trend over 2001–2020 (natural year-to-year variability dominates), while annual minimum temperature shows a statistically significant warming trend of approximately $+0.91^\circ\text{C}$ over the same period — a finding with direct implications for night-time heat stress and cold-storage planning, independent of any change in rainfall.

The 2021–2025 forecast generated by this pipeline was subsequently validated against genuinely observed data for the same period, and this post-hoc comparison strongly reinforces the project's original conclusions rather than contradicting them. At the aggregation levels that matter most for planning — season-wise totals ($R^2 = 0.83$) and month-of-year climatology ($R^2 = 0.79$) — the model's predictions align closely with what actually occurred, correctly identifying the Northeast Monsoon as the dominant rainfall period and November as the wettest month in every comparison. Rain-presence (wet/dry) classification against actual data improved from 65.1% at the day level to 91.7% at the month level and a full 100% at the year level, confirming that the model's broader seasonal judgment is dependable even when exact daily timing is not. As anticipated in the original limitations section, accuracy declines sharply at the year-

wise ($R^2 = -1.05$) and day-wise ($R^2 = -0.10$) levels, since single-day and single-year rainfall totals are governed by short-lived synoptic weather systems that a station-level statistical model cannot resolve — a limitation of this class of model generally, not a defect specific to this pipeline, and one that has now been demonstrated on real, independent data rather than only inferred from the test set.

Taken together, the original analysis and the subsequent validation converge on the same practical conclusion: this pipeline is a reliable tool for seasonal and climatological planning — irrigation scheduling, crop-cycle advisories, and early-warning preparation ahead of the Northeast Monsoon — and should be used alongside, not in place of, official IMD short-range forecasts for day-to-day operational decisions. The project therefore fulfils its stated objectives of characterising Karaikal's monsoon pattern, building an evidence-based forecasting model, and providing a transparently validated, data-driven foundation for the adaptation strategies recommended in this report.