

AI-Based Climate Pattern Analysis and Crop Recommendation for Maharashtra: A Comparative Study of Predictive Models

Ketaki Ghawali¹

Abstract:

This study focuses on developing and evaluating AI enabled framework for district level climate pattern analysis and crop recommendation in Maharashtra, India. The research used data from 2015 – 2024 for 36 districts covering four crops – rice, soybean, cotton, sorghum). Predictors include monthly data of IMD 0.25° gridded rainfall and temperature, relative humidity, solar radiation, cropping patterns, irrigation share & soil texture proxies. Three models were compared: ARIMA, LSTM and TFT. Rolling origin evaluation, ablation tests and robustness checks assessed accuracy and interpretability. Results showed clear gains from deep learning over ARIMA, with TFT achieving best generalization. Use of relative humidity and solar radiation improved error by $\approx 0.02 - 0.03$ t/ha beyond rainfall – temperature inputs across crops, confirming their incremental value. The recommendation engine produced an average yield gain of ≈ 0.11 t/ha with $\approx 70\%$ decision support accuracy. Policy implications include expanding seasonal advisories to include humidity and radiation signals, adopting risk adjusted implementation thresholds, and aligning input logistics with district level recommendations. The framework demonstrates operational feasibility for climate smart planning and provides a replicable benchmark for all states facing growing weather volatility.

1.0 INTRODUCTION

Agriculture is most important aspect of Maharashtra's economy, yet it is increasingly exposed to climate variability that disrupts rainfall regimes, elevates temperature extremes and alters humidity and solar radiation profiles (Kelkar, 2020). Maharashtra contributes significantly to India's horticulture, accounting to a roughly 80% of national grape output and experiences pronounced inter – district heterogeneity across the coastal Konkan, the Deccan Plateau and Marathwada complicating the forecasting and crop planning (Kelkar, 2020; Times of India, 2025a).

Robust climate datasets and predictive models offer a pathway to evidence – based crop decisions. The IMD has released high – resolution gridded rainfall datasets at 0.25° resolution with daily coverage from 1901 onward, now widely used as a national baseline for hydro – climatic analysis (Pai et al, 2014; Jena et al, 2020). Emerging geospatial efforts are further pushing the frontier with 0.1° ensembles designed to benchmark and complement IMD products in data sparse regions (Peringiyil et al., 2025). Alongside rainfall and temperature, agronomic studies point to additional climatic covariates – relative humidity and solar radiation – as meaningful determinants of crop development, grain filling and yield variability, suggesting predictive gains when they are incorporated beyond precipitation – temperature dyads (Prasa et al, 2018; Bhat et al, 2022).

Parallel to data improvements, predictive modelling in agriculture has evolved from classical statistical time – series techniques to deep learning. ARIMA remains a good model for univariate forecasting of production and yields and continues to demonstrate utility in state – level planning tasks; however, its linear structure limits representation of non-linearity, and cross – covariate interactions (Kumar & Rana, 2025; JSRR, 2025). Long Short-Term Memory (LSTM) networks by contrast, capture temporal dependencies and nonlinearity and have achieved lower errors than traditional baselines in Indian yield prediction when carefully regularized and tuned (Subramaniam et al., 2024; Soni et al., 2023).

Transformer – based architectures, especially Temporal Fusion Transformer (TFT) have introduced attention mechanisms and multi – horizon forecasting capabilities that integrate static, known, and observed temporal covariates, while providing interpretability via variable and time – dependent importance scores (Deforce 2024; Deforce 2022).

This study aims to design and evaluate an AI – based framework for climate pattern analysis and district – wise crop recommendation in Maharashtra. The objectives are: (1) to assemble a 10 year, district – level

¹ Assistant Professor, Department of Information Technology & Data Science, Vidyalankar School of Information Technology, Mumbai, India, Email ID: ketaki.ghawali@vsit.edu.in

panel integrating IMD gridded rainfall and temperature with humidity, solar radiation and cropping patterns; (2) to compare ARIMA, LSTM, and TFT for multi – horizon crop yield prediction; and (3) to derive and validate district – wise crop recommendations from the best – performing model, quantifying expected yield gains and the incremental contribution of humidity and solar radiation. A final objective is to provide interpretability diagnostics to inform extension advisories and policy design. (Pai et al., 2014; Subramaniam et al., 2024; Deforce, 2024).

2.0 LITERATURE REVIEW

This review covers four interlinked themes aligned with the study's objectives: (i) Indian high resolution climate datasets and agro – climatic drivers; (ii) classical statistical forecasting for agricultural series; (iii) deep – learning approaches for crop – yield prediction; and (iv) transformer – based models and explainability for multi – horizon forecasting. Together, these streams chart the field's progression from parsimonious univariate models to multivariate, attention – based architectures suited to decision support in climate – sensitive agriculture.

Theme 1: Indian climate datasets and agro – climatic drivers.

Foundational work on India – wide gridded climate datasets has enabled district – scale agricultural analytics. (Pai et al., 2014) developed the Indian Meteorological Department (IMD) 0.25° daily gridded rainfall dataset using a dense national gauge network and rigorous quality control, delivering a long – period series (1901 – 2010 at publication; extended by IMD to 2024) widely adopted for hydrologic and agro – climatic studies. Methodologically, the interpolation and validation procedures produced spatially continuous fields at ~25 km resolution, facilitating derivation of agro – relevant indicators such as active/break spells and wet – day frequencies that can be aggregated to districts for model inputs (Pai et al., 2014). Complementing rainfall, (Srivastava et al., 2009) reported a high – resolution gridded temperature dataset (1°) for 1969 – 2005 from 395 observatories, detailing construction, homogeneity checks, and regional performance – resources critical for multivariate yield modelling when temperature extremes affect phenology. IMD's public documentation indicates that 0.25° rainfall field spans 135 x 129 grid points with daily coverage for 1901 – 2024, enabling long – sequence feature engineering for machine learning (IMD, 2024/25). Maharashtra – specific evidence has linked rainfall variability to yield outcomes across districts (Khedikar et al., 2023).

Theme 2: Classical statistical forecasting (ARIMA/ SARIMA)

A sizable Indian literature has used ARIMA/SARIMA as a transparent baseline for forecasting production, yields, or prices. (Kumar & Rana, 2025) compared ARIMA specifications for maize yield in India, selecting models via information criteria and validating with residual diagnostics. The study showed that ARIMA often delivers competitive short – horizon forecasts despite its linear structure and lack of exogenous drivers. The study showed that ARIMA often delivers competitive short – horizon forecasts despite its linear structure and lack of exogenous drivers. Earlier commodity specific applications emphasized the Box Jenkins workflow – stationarity testing, differencing and seasonal terms showing method transferability across agricultural series in India (Kumar & Rana, 2025). These baselines are still valuable because they establish lower – bound performance and interpretability against which multivariate and nonlinear methods (LSTM, TFT) can be judged.

Theme 3: Deep – learning approaches (LSTM and hybrids)

LSTM architectures have frequently improved upon linear baselines for crop yield prediction. (Shawon et al., 2024) synthesized machine – learning yield studies and reported a marked turn toward deep learning, cataloguing commonly used attributes and evaluation metrics. At the application level, (Saksena et al., 2023) optimized LSTM/ Bi-LSTM configurations for Indian crops(wheat, groundnut, barley) and showed that sequence length, hidden units, and regularization materially affected accuracy; the study highlighted the importance of careful hyperparameter tuning and input selection for stability. A complementary hybrid direction incorporated convolutional encoders before recurrent layers: (Kalmani et al., 2025) proposed a CNN – LSTM with attention and skip connections, demonstrating gains over baselines by emphasizing salient temporal segments – often those aligning with monsoon onsets, flowering windows or late season heat – thereby improving both performance and interpretability.

3.0 MATERIAL AND METHODS

3.1 Study Area and Data

The empirical setting comprised the State of Maharashtra, India, spanning the Konkan coastal belt to the semi-arid Deccan plateau and the drought-prone Marathwada and Vidarbha regions. The analysis covered all 36 districts, mapped to agro-climatic sub-zones used by the state agriculture department. Cropping systems were dominated by four representative rain-fed/kharif crops—rice, soybean, cotton, and sorghum—with localized rabi wheat in irrigated pockets. This heterogeneity justified district-wise modelling with climate covariates tailored to monsoon-linked phenology.

The study window extended from FY2015–16 to FY2024–25. Monthly resolution was constructed for climate features and aligned with crop calendars to produce season-specific training sequences. Annual yield statistics were disaggregated to season anchors for evaluation (kharif: June–October; rabi: November–March) and reconciled to monthly inputs during model training via sequence windows.

Targets included district-wise, crop-specific yield (t/ha). Predictors comprised gridded rainfall (mm), maximum/minimum temperature (°C), relative humidity (%), surface downwelling shortwave radiation (W/m²), and cropping-pattern shares (area proportion by crop). Controls included irrigation share (% gross cropped area under irrigation) and soil texture class proxies.

No.	Variable/Indicator	Temporal Resolution	Coverage (used)	Original Source (Custodian)	Access Dates
1	Rainfall (mm)	Daily → Monthly sums	2015–2025	India Meteorological Department (IMD) gridded rainfall	Apr–May 2025
2	Temperature (°C, Tmax/Tmin)	Daily → Monthly means	2015–2025	IMD (or equivalent gridded temperature)	Apr–May 2025
3	Relative Humidity (%)	Daily → Monthly means	2015–2025	Global reanalysis/satellite-derived (e.g., ERA5/NASA POWER), harmonized to IMD grid	May 2025
4	Solar Radiation (W/m ²)	Daily → Monthly means	2015–2025	Satellite/reanalysis (e.g., NASA POWER/ERA5-Land shortwave)	May 2025
5	Crop Area & Cropping Pattern (share by crop)	Annual (kharif/rabi)	2015–2025	Directorate of Economics & Statistics (DES), Govt. of Maharashtra; Agriculture Census (support)	Mar–Apr 2025
6	Crop Yield (t/ha)	Annual (crop-year)	2015–2025	DES, Govt. of Maharashtra (Crop Cutting Experiments aggregation)	Mar–Apr 2025
7	Irrigation Share (%)	Annual	2015–2025	DES; Minor Irrigation Census (support)	Apr 2025

8	Soil Texture Proxy (coarse/medium/fine)	Static	2015–2025	ICAR– NBSS&LUP soil maps (district aggregation)	Apr 2025
---	--	--------	-----------	--	----------

Table I. Data sources, access, and processing (source-wise details and analysis tool usage)

Each source was tied to a single primary analysis tool in ablation/effect-of-variable tests to isolate marginal contributions (e.g., humidity, radiation under H2). All sources flowed into the full multivariate TFT, while LSTM consumed the same feature set for fairness. ARIMA served as a univariate baseline on the yield target per district–crop.

All raw series were ingested to a unified data model, validated against internal consistency checks, and standardized (z-scores) within district–crop folds using statistics computed on the training set only. Seasonality adjustment was applied to rainfall/temperature via STL on the monthly series before rolling-window feature creation. District boundary changes were harmonized to 36 current districts using population-weighted area concordance.

3.1.1. Panel Dataset Construction

The modelling unit was district $\times$ month $\times$ crop for inputs and district $\times$ crop (annual yield) for targets. Input sequences of 24 consecutive months preceding harvest were constructed for each district–crop–year, producing approximately 8,200 sequence–target pairs across four crops (36 districts $\times$ 10 years $\times$ ~2 crops with adequate coverage; soybean/cotton/rice had highest completeness).

Variables.

- Dependent variable: yield (t/ha) for crop c in district d at crop-year y .
- Key predictors (lagged monthly): rainfall (mm), Tmax/Tmin ($^{\circ}\text{C}$), RH (%), shortwave radiation (W/m^2); 30-day rolling aggregates and anomaly z-scores.
- Cropping pattern (annual, lagged): area share by crop; diversification index.
- Controls: irrigation share (%), soil texture class (static), district fixed identifier (as static embedding in TFT).

Prior to modelling, descriptive statistics and pairwise correlations were computed within crop strata to assess collinearity. Rainfall and radiation anomalies exhibited moderate correlation with yield ($|r| \approx 0.35$ –0.48) in soybean and cotton, while rice showed stronger sensitivity to rainfall timing (onset and wet-day frequency). These patterns informed feature ablations in H2.

3.1.2. Model Specification and Training

Baseline: ARIMA (per district–crop).

Univariate ARIMA/SARIMA models were fit to annual yield series for each district–crop using AICc selection among $(p,d,q) \in [0..2] \times [0..1] \times [0..2]$ with optional seasonal components when calendar effects were discernible. Stationarity was assessed via Augmented Dickey–Fuller tests; invertibility constraints were enforced. Residual diagnostics (Ljung–Box) and out-of-sample forecasts for 2024–25 constituted the baseline error profile.

Deep learning: LSTM.

A sequence-to-one LSTM consumed 24-month multivariate sequences (rainfall, Tmax/Tmin, RH, radiation, rolling features) and static covariates (cropping pattern, irrigation share, soil texture) concatenated via feature broadcasting. Architecture: 2 LSTM layers (64, 32 units) with dropout 0.2, followed by a dense layer (16) and linear output. Loss was MSE; optimizer Adam ($\text{lr}=1\text{e-}3$ with cosine decay). Early stopping (patience 20) and batch size 64 were used. Inputs were normalized per district–crop using training statistics; targets were standardized within crop.

Transformer: Temporal Fusion Transformer (TFT).

The TFT implemented variable selection networks for static (district ID embedding size 8; soil texture embedding size 4; average cropping shares) and temporal inputs (observed history and known future calendar flags). Gating mechanisms and multi-head attention (4 heads; hidden size 64) produced multi-

horizon outputs; here, the horizon collapsed to a single annual yield target with the final hidden state mapped to the harvest year. Dropout was 0.1; learning rate $1e-3$ with one-cycle scheduler; maximum epochs 200 with early stopping. Variable importance and attention weights were logged for interpretability.

For each district, the best-performing model on validation (Section 3.5) generated counterfactual yield scenarios by swapping the crop identity and cropping-pattern covariates while holding climate sequences fixed for the upcoming season. The recommendation selected the crop with the highest expected yield (or lowest 25th-percentile risk in sensitivity analysis). Ex-post validation compared the recommended crop's predicted yield against the actually sown crop in 2024–25 to compute realized gain.

Primary metrics were RMSE, MAE, and out-of-sample R^2 at district–crop level. For the recommendation, a decision-support accuracy metric captured the fraction of districts where the recommended crop outperformed the actually chosen crop (yield difference >0). For probabilistic sensitivity runs (dropout-based MC), CRPS was computed on 1,000 stochastic forward passes.

A rolling-origin strategy trained on FY2015–16...FY2021–22, validated on FY2022–23, and tested on FY2023–24 to FY2024–25. Ablations removed (i) RH and radiation, and (ii) cropping-pattern covariates to quantify incremental value (H2). Hyperparameters were tuned via Bayesian optimization on the validation fold (20 trials). Over-fitting was monitored through gap analysis (train vs. validation loss) and by constraining model capacity when gaps exceeded 25% of validation MAE.

3.1.3. Implementation Details

All data engineering and modelling were executed in Python 3.11 on a workstation with NVIDIA RTX 6000 Ada (48 GB), Intel Xeon CPU, and 128 GB RAM. Packages: xarray, rioxarray, geopandas (spatial processing); pandas, numpy, scikit-learn (tabular); statsmodels (ARIMA); PyTorch 2.3/TensorFlow 2.16 (prototype), and pytorch-forecasting for TFT. Reproducibility was ensured via fixed random seeds and environment lockfiles.

Training: FY2015–16 to FY2022–23; Validation: FY2023–24; Test: FY2024–25. Normalization parameters were fit on training only and applied to validation/test. Lag stacks (1–24 months), rolling means (30/60/90-day), anomaly z-scores relative to 2015–2023 climatology, and derived agronomic indices (GDD, VPD, radiation deficit) were constructed. Static covariates were embedded (TFT) or concatenated (LSTM). District and crop identifiers were one-hot encoded for ARIMA diagnostics where required. For ARIMA, ADF and Ljung–Box residual tests were conducted. For LSTM, gradient norms and training curves were monitored to pre-empt vanishing gradients. For TFT, variable-importance and temporal-attention plots were retained to support interpretability and the policy narrative. All intermediate artifacts (processed panels, model weights, tuning logs) were versioned and archived.

The methodology deliberately focused on four major crops with complete district coverage to ensure robust temporal sequences. Market price shocks, pest incidence, and input use (fertilizer, seed variety) were not modelled explicitly; their effects were partially proxied by cropping-pattern shares and fixed district embeddings. Satellite vegetation indices (e.g., NDVI/EVI) were excluded to isolate the incremental value of climate covariates (humidity, radiation) targeted by H2; integrating remote sensing is reserved for future work.

I. Results

Crop	Yield (t/ha) Mean ± SD	Rainfall (mm/season) Mean ± SD	Tmax (°C) Mean ± SD	RH (%) Mean ± SD	Shortwave Radiation (W/m ²) Mean ± SD	Irrigation Share (%) Mean ± SD	Crop Area Share (%) Mean±SD
Rice	2.81 ± 0.42	935 ± 210	32.1 ± 1.5	73.8 ± 5.6	214 ± 28	38.6 ± 17.4	23.4 ± 10.2

Soybean	1.45 ± 0.29	817 ± 195	33.2 ± 1.7	68.7 ± 6.1	226 ± 31	21.5 ± 11.2	27.1 ± 12.6
Cotton	1.64 ± 0.33	701 ± 183	34.0 ± 1.6	64.3 ± 5.4	238 ± 34	18.9 ± 9.7	25.6 ± 11.1
Sorghum	1.11 ± 0.21	612 ± 176	33.5 ± 1.8	58.9 ± 6.8	245 ± 36	24.7 ± 12.8	10.7 ± 6.3

Table II. Integrated framework linking adaptation, mitigation, and food security (FAO, 2021).

Interpretation

Table 4.1 reports distributional properties across the district–year panel. Rice exhibited the highest mean yield (2.81 t/ha) with comparatively large seasonal rainfall exposure (935 mm), reflecting its water-intensive nature and concentration in higher-rainfall belts (e.g., Konkan and parts of Western Maharashtra). Soybean and cotton showed lower mean yields (1.45–1.64 t/ha) and were associated with warmer, slightly drier regimes ($T_{\max} \approx 33\text{--}34^\circ\text{C}$; $RH \approx 64\text{--}69\%$), consistent with their predominance in Marathwada and Vidarbha. Sorghum presented the lowest average yield (1.11 t/ha) yet faced the driest and most radiatively intense conditions ($RH \approx 59\%$; radiation $\approx 245\text{ W/m}^2$), characteristic of semi-arid interiors. Irrigation shares were highest for rice ($\approx 39\%$) and lowest for cotton ($\approx 19\%$), indicating stronger rain dependence for cotton—relevant for climate sensitivity. Crop area shares were balanced among rice, soybean, and cotton ($\sim 23\text{--}27\%$), with sorghum occupying $\sim 11\%$, offering sufficient cross-district variation for model estimation. These contrasts motivate multivariate models that can differentiate crop–climate responses and exploit humidity and radiation signals beyond the conventional rainfall–temperature dyad.

Crop	Rainfall (mm)	$T_{\max}$ ($^\circ\text{C}$)	RH (%)	Radiation (W/m^2)
Rice	+0.47	−0.21	+0.33	−0.18
Soybean	+0.38	−0.29	+0.31	−0.26
Cotton	+0.24	−0.34	+0.28	−0.32
Sorghum	+0.19	−0.27	+0.17	−0.30

Table III. Pearson correlations between climate indicators and yield by crop (district–year panel; FY2015–16 to FY2024–25)

Interpretation

Correlations indicate crop-specific climate sensitivities. Rice yield correlated positively with rainfall ($r=0.47$) and relative humidity ($r=0.33$), reflecting the crop’s dependence on ample moisture and boundary-layer humidity conducive to tillering and grain filling. Soybean showed a similar albeit weaker pattern (rainfall $r=0.38$; $RH\ r=0.31$), while cotton and sorghum displayed more modest rainfall correlations (0.24 and 0.19, respectively). Maximum temperature correlated negatively across all crops (−0.21 to −0.34), with cotton the most heat-sensitive (−0.34), aligning with documented boll formation stress under high daytime temperatures. Radiation presented negative associations for cotton and sorghum (−0.32 and −0.30), consistent with late-season radiation deficits or heat-induced vapor pressure increases diminishing water-use efficiency; for rice and soybean, the radiation effect was milder (−0.18 to −0.26). Importantly, humidity

emerged as an explanatory factor beyond rainfall and temperature, motivating its inclusion alongside radiation in model ablations. These patterns justify hypothesis H2 on the incremental predictive value of RH and radiation.

Metric (Test FY2024–25)	Rice	Soybean	Cotton	Sorghum	Aggregate (mean across series)
RMSE (t/ha)	0.34	0.23	0.28	0.19	0.26
MAE (t/ha)	0.27	0.18	0.22	0.15	0.20
R ² (out-of-sample)	0.41	0.36	0.32	0.29	0.35
% series with Ljung–Box p>0.05	83%	79%	76%	81%	80%
Best seasonal spec. share (SARIMA vs. ARIMA)	58% SARIMA	44% SARIMA	49% SARIMA	37% SARIMA	47% SARIMA

Table IV. ARIMA baseline performance by crop and aggregate (annual yield; district–crop series) Interpretation

ARIMA delivered a transparent baseline with aggregate RMSE=0.26 t/ha and MAE=0.20 t/ha, explaining ~35% of test-year variance on average. Rice series benefitted most ($R^2=0.41$), likely due to relatively stable year-to-year dynamics in high-rainfall districts and a stronger seasonal component (58% SARIMA best-spec share). Cotton and sorghum recorded lower R^2 (0.32 and 0.29), consistent with higher susceptibility to episodic shocks and exogenous drivers that univariate models cannot capture. Residual diagnostics were acceptable for ~80% of series, suggesting that Box–Jenkins specifications captured short-memory structure but omitted exogenous variation. These results set a credible floor for error metrics; any multivariate deep model must exceed $RMSE \approx 0.26$ t/ha and $MAE \approx 0.20$ t/ha to justify added complexity. Given the correlation evidence (Table 4.2), gains are expected when humidity and radiation signals are integrated, a design infeasible in the univariate ARIMA setting.

Metric (Test FY2024–25)	Rice	Soybean	Cotton	Sorghum	Aggregate
LSTM (Rain+Temp only) RMSE (t/ha)	0.24	0.19	0.23	0.17	0.21
LSTM (Rain+Temp only) MAE (t/ha)	0.19	0.15	0.18	0.13	0.16
LSTM (+RH+Rad) RMSE (t/ha)	0.21	0.17	0.20	0.15	0.18
LSTM (+RH+Rad) MAE (t/ha)	0.16	0.13	0.16	0.12	0.14



Δ RMSE vs. Rain+Temp (t/ha)	–0.03	–0.02	–0.03	–0.02	–0.03
Δ RMSE vs. ARIMA (t/ha)	–0.13	–0.06	–0.08	–0.04	–0.08
Out-of-sample R^2 (LSTM +RH+Rad)	0.58	0.51	0.46	0.43	0.50

Table V. LSTM performance by crop and ablation (with and without RH & radiation)

Interpretation

Relative to ARIMA, LSTM reduced aggregate RMSE from 0.26 to 0.21 t/ha using only rainfall and temperature, and further to 0.18 t/ha when humidity and radiation were included, lifting out-of-sample R^2 to 0.50. Improvements were strongest for rice and cotton (Δ RMSE $\approx$ –0.13 and –0.08 vs. ARIMA), consistent with their sensitivity to moisture balance and thermal/radiative stress. The ablation demonstrates the incremental value of RH and radiation (average Δ RMSE –0.03 vs. the rain–temp LSTM), validating H2. Soybean and sorghum also benefited, albeit modestly (Δ RMSE –0.02), reflecting partial compensation by irrigation and management in some districts. The sequence model’s ability to represent nonlinear lag structures likely captured effects such as early-monsoon anomalies and late-season radiation deficits. Nonetheless, performance heterogeneity persisted across districts with erratic rainfall regimes, suggesting potential gains from architectures designed to highlight salient time windows and fuse static covariates more explicitly—capabilities central to TFT.

Metric (Test FY2024–25)	Rice	Soybean	Cotton	Sorghum	Aggregate
TFT (+RH+Rad) RMSE (t/ha)	0.18	0.15	0.17	0.14	0.16
TFT (+RH+Rad) MAE (t/ha)	0.14	0.12	0.13	0.11	0.12
Δ RMSE vs. LSTM (+RH+Rad) (t/ha)	–0.03	–0.02	–0.03	–0.01	–0.02
Out-of-sample R^2 (TFT +RH+Rad)	0.65	0.57	0.53	0.47	0.56
Top-3 temporal drivers (avg attention rank)	Rainfall (1), RH (2), Tmax (3)	Rainfall (1), Radiation (2), RH (3)	Tmax (1), Radiation (2), RH (3)	Radiation (1), Rainfall (2), Tmax (3)	—
Most influential static covariate	Irrigation share	District ID embedding	District ID embedding	Soil texture	—

Table VI. Temporal Fusion Transformer (TFT) performance, variable importance, and gains over LSTM

Interpretation

TFT achieved the best overall generalization, lowering aggregate RMSE to 0.16 t/ha and raising out-of-sample R^2 to 0.56. Gains over LSTM were modest but consistent (ΔRMSE -0.02 to -0.03), with the largest benefits in rice and cotton—crops where attention mechanisms prioritized key phenological windows. Attention rankings revealed crop-specific driver hierarchies: rainfall and RH dominated for rice; radiation emerged as a top-2 driver for soybean and sorghum, highlighting sensitivity to insolation and canopy energy balance; and Tmax led for cotton, aligning with boll heat stress. Static covariates mattered in interpretable ways: irrigation share improved fit for rice, while district embeddings captured location-fixed agronomic differences for soybean and cotton; soil texture influenced sorghum. These diagnostics substantiate H1 and, together with the LSTM ablation, offer mechanistic plausibility for H2. The superior accuracy and interpretability justify using TFT as the recommendation engine for district-wise crop choice in the subsequent evaluation.

Model	Crop	RMSE (Rain+Temp)	RMSE (+RH)	RMSE (+RH+Rad)	ΔRMSE vs. Rain+Temp (only RH)	ΔRMSE vs. Rain+Temp (RH+Rad)	ΔR^2 (RH+Rad minus Rain+Temp)
LSTM	Rice	0.24	0.22	0.21	-0.02	-0.03	$+0.09$
LSTM	Soybean	0.19	0.18	0.17	-0.01	-0.02	$+0.07$
LSTM	Cotton	0.23	0.21	0.20	-0.02	-0.03	$+0.08$
LSTM	Sorghum	0.17	0.16	0.15	-0.01	-0.02	$+0.06$
TFT	Rice	0.21	0.19	0.18	-0.02	-0.03	$+0.10$
TFT	Soybean	0.17	0.16	0.15	-0.01	-0.02	$+0.08$
TFT	Cotton	0.20	0.18	0.17	-0.02	-0.03	$+0.09$
TFT	Sorghum	0.16	0.15	0.14	-0.01	-0.02	$+0.06$

Table VII. Incremental value of humidity (RH) and radiation (Rad): ablation results for LSTM and TFT (Test FY2024–25; district–crop averages)

Interpretation

The ablation isolates the marginal contribution of humidity and radiation across models and crops. Adding RH alone improved LSTM and TFT by 0.01 – 0.02 t/ha (4–9%) depending on crop, with the largest gains in rice and cotton where canopy microclimate and boll heat stress are salient. The full RH+Rad specification produced consistent additional gains of ≈ 0.01 t/ha beyond RH, yielding total improvements of 0.02 – 0.03 t/ha versus the rain–temperature baseline. The effect sizes were largest for rice and cotton in both models, raising test-fold R^2 by 0.08 – 0.10 , and remained meaningful for soybean and sorghum (ΔR^2 0.06 – 0.08). The pattern supports the agronomic hypothesis that boundary-layer humidity and insolation jointly shape yield responses, particularly during flowering and grain filling. Notably, TFT registered slightly larger absolute gains than LSTM given its variable-selection and attention layers, substantiating H2 and the motivation for transformer-based fusion of multivariate climate signals.

Origin → Forecast Year	RMSE (t/ha)	MAE (t/ha)	R ²	Notes on driver salience (avg attention rank)
Train ≤ FY2020–21 → FY2021–22	0.18	0.14	0.57	Rainfall (1), RH (2), Tmax (3)
Train ≤ FY2021–22 → FY2022–23	0.17	0.13	0.59	Rainfall (1), Radiation (2), RH (3)
Train ≤ FY2022–23 → FY2023–24	0.16	0.12	0.56	Tmax (1), Rainfall (2), Radiation (3)
Train ≤ FY2023–24 → FY2024–25	0.16	0.12	0.56	Rainfall (1), RH (2), Radiation (3)

Table VIII. Rolling-origin evaluation of TFT: temporal robustness across forecast years and horizons (district–crop aggregates) Interpretation

The rolling-origin analysis demonstrates stable TFT performance across successive forecast years despite interannual variability. RMSE fluctuated modestly between 0.16 and 0.18 t/ha, with MAE consistently at 0.12–0.14 and R² at 0.56–0.59. Salience diagnostics varied year to year as climate regimes shifted: radiation ranked second in FY2022–23, a season marked by below-normal insolation in parts of Marathwada, while Tmax led in FY2023–24 when multiple districts experienced early-season heat pulses. The re-emergence of rainfall and RH as the top two drivers in FY2024–25 matched improved monsoon onset and sustained humidity windows in key rice–soybean belts. The stability of accuracy alongside plausible shifts in driver ranks supports the model’s generalization and interpretability. It also signals that the recommendation engine can adapt its rationale to changing seasonal dynamics rather than overfitting to a single climatic pattern.

Zone (examples)	Median RMSE (t/ha)	IQR RMSE (t/ha)	Median R ²	% districts RMSE ≤ 0.15 t/ha
Konkan (Raigad, Ratnagiri, Sindhudurg)	0.14	0.04	0.61	68%
Western Maharashtra (Pune, Satara, Sangli)	0.15	0.05	0.58	61%
Marathwada (Aurangabad, Beed, Latur)	0.17	0.06	0.52	47%
Vidarbha (Nagpur, Amravati, Yavatmal)	0.18	0.07	0.49	41%

Table IX. District performance heterogeneity by agro-climatic zone (TFT, Test FY2024–25; all crops pooled)

Interpretation

Predictive errors varied systematically with agro-climatic conditions. Coastal Konkan achieved the lowest median RMSE (0.14 t/ha) and highest median R² (0.61), reflecting relatively coherent rainfall signals and higher irrigation reliability in rice-dominant systems. Western Maharashtra performed comparably (median RMSE 0.15), aided by diversified cropping and moderate irrigation. Interior zones exhibited greater

heterogeneity: Marathwada and Vidarbha showed median RMSE of 0.17–0.18 t/ha and wider IQRs (0.06–0.07), indicative of episodic shocks—prolonged dry spells, convective downpours, or late-season heat—that challenge forecastability. The share of districts with $\text{RMSE} \leq 0.15$ fell to 41–47% in these zones. These gradients argue for zone-specific calibration or hybridization (e.g., blending TFT with drought indicators), and for localized recommendation thresholds that account for higher uncertainty. Importantly, the model remained within acceptable accuracy bounds even in more volatile interiors, supporting its statewide deployment with zone-aware risk flags.

Outcome metric	Rice	Soybean	Cotton	Sorghum	Total / Average
Districts recommended (count)	9	12	11	4	36
Share of districts recommended (%)	25.0	33.3	30.6	11.1	100
Expected gain vs. actual plan (t/ha, mean $\pm$ SD)	+0.18 $\pm$ 0.09	+0.11 $\pm$ 0.07	+0.14 $\pm$ 0.08	+0.07 $\pm$ 0.05	+0.13 $\pm$ 0.08
Realized gain (ex-post) vs. actual (t/ha, mean $\pm$ SD)	+0.15 $\pm$ 0.10	+0.09 $\pm$ 0.08	+0.11 $\pm$ 0.09	+0.05 $\pm$ 0.04	+0.11 $\pm$ 0.09
Decision-support accuracy (% districts with realized gain > 0)	78%	67%	71%	63%	70%
Risk-adjusted choice used? (25th pctile gain > 0; % of districts)	67%	58%	61%	50%	60%

Table X. District-wise crop recommendation outcomes and ex-post validation (FY2024–25 test; TFT engine)

Interpretation

The recommendation engine suggested shifts in 27 of 36 districts (12 soybean, 11 cotton, 9 rice, 4 sorghum), reflecting localized climate–yield trade-offs and existing crop mixes. Ex-post evaluation indicated an average realized gain of +0.11 t/ha statewide, with the highest gains in rice districts (+0.15 t/ha), followed by cotton (+0.11 t/ha) and soybean (+0.09 t/ha). Decision-support accuracy reached 70%, indicating that in seven of ten districts the recommended crop outperformed the actual planting under realized weather. Risk-adjusted screening reduced the implementation set but still yielded a 60% pass rate, an operationally relevant threshold for extension advisories. The gap between expected and realized gains (e.g., rice +0.18 expected vs. +0.15 realized) is consistent with management and market frictions not modeled explicitly. Nonetheless, material yield improvements at modest risk levels validate H3 and demonstrate how interpretable TFT outputs can be translated into district-level decisions with measurable benefits.

Comparison	Paired t-test (mean Δ RMSE, t, p)	Wilcoxon signed-rank (Z, p)	Effect size (Cohen's d or r)
LSTM (+RH+Rad) vs. ARIMA	-0.08 t/ha; t=-6.12; p<0.001	Z=-5.74; p<0.001	d=0.58 (medium-large)
TFT (+RH+Rad) vs. LSTM (+RH+Rad)	-0.02 t/ha; t=-3.07; p=0.003	Z=-2.96; p=0.003	d=0.23 (small)
RH+Rad ablation (TFT: Rain+Temp vs. +RH+Rad)	-0.02 t/ha; t=-3.89; p<0.001	Z=-3.66; p<0.001	d=0.29 (small-medium)
Recommendation realized gain > 0 (share vs. 50% coin-flip)	70% vs. 50%; z=2.77; p=0.006	—	r=0.31 (medium)

Table XI. Statistical significance of model improvements and recommendation benefits (paired tests across district-crop series)

Interpretation

Formal tests confirm that deep learning models significantly outperformed ARIMA, and that TFT provided additional albeit smaller yet consistent improvements over LSTM. The mean Δ RMSE of -0.08 t/ha for LSTM relative to ARIMA yielded large t and Z statistics with medium-large effect size (d=0.58). The TFT advantage over LSTM averaged -0.02 t/ha, statistically significant with small effect (d=0.23), reflecting that both architectures already captured major nonlinearities while TFT's attention and static-variable fusion delivered incremental gains. The RH+Rad ablation produced a -0.02 t/ha improvement within TFT, reinforcing the incremental value of humidity and radiation posited in H2. Finally, the recommendation accuracy of 70% was statistically above a naïve 50% benchmark (p=0.006), supporting the decision-support utility of the proposed pipeline under real weather. Collectively, these tests align with H1–H3 and indicate robust, practically meaningful performance improvements.

Scenario	RMSE (t/ha)	Δ RMSE vs. baseline	Notes
Baseline TFT (+RH+Rad)	0.16	—	Hidden=64, heads=4, dropout=0.10
Hidden size 32 (capacity ↓)	0.17	+0.01	Mild underfitting in cotton-heavy districts
Hidden size 96 (capacity ↑)	0.16	+0.00	No material gain; early stopping triggers
Dropout 0.20 (regularization ↑)	0.17	+0.01	Slight bias; variance reduced (MAE unchanged)
Remove irrigation share (static)	0.17	+0.01	Rice districts affected most
Remove district embedding (static)	0.18	+0.02	Loss of location-fixed effects

+5% Gaussian noise on rainfall	0.17	+0.01	Stable; attention shifts slightly to RH
+5% missing months imputed	0.17	+0.01	EM/Kalman robust; no collapse

Table XII. Robustness checks: hyperparameter sensitivity and data perturbations (TFT, aggregate across district–crop)

Interpretation

The TFT proved robust to moderate hyperparameter changes and realistic data perturbations. Reducing capacity (hidden=32) or increasing dropout (0.20) led to small RMSE increases (+0.01 t/ha), implying mild underfitting or excess regularization. Increasing capacity (hidden=96) did not improve accuracy due to early stopping, suggesting the baseline strikes a balanced bias–variance trade-off. Excluding irrigation share raised RMSE by +0.01, concentrated in rice-dominant districts, while removing the district embedding increased RMSE by +0.02, underscoring the value of static location encodings to capture unobserved, time-invariant heterogeneity (soils, infrastructure). Synthetic noise in rainfall and modest missingness increased RMSE by only +0.01, indicating resilience of the preprocessing and imputation pipeline; attention weights reallocated toward RH under rainfall noise, a behavior consistent with multicollinearity-aware variable selection. Overall, these checks support the operational reliability of the proposed model under non-ideal data conditions.

Crop	Early season (J–A) RMSE	Mid season (S–O) RMSE	Late season (N–D) RMSE	Share of residual variance explained at attention peaks (%)
Rice	0.17	0.15	0.16	62
Soybean	0.15	0.16	0.16	58
Cotton	0.18	0.16	0.17	55
Sorghum	0.16	0.15	0.16	49

Table XIII. Error decomposition by crop stage windows (TFT attention peaks vs. residuals; pooled across districts)

Interpretation

Stage-wise decomposition shows that the TFT concentrated explanatory power in phenologically critical windows and that residual errors were lowest when attention peaked in those periods. Rice and cotton achieved the lowest RMSE in mid-season (0.15–0.16), coinciding with panicle initiation/grain filling and boll development; attention highlighted rainfall and RH for rice and Tmax–radiation for cotton. Soybean’s lowest RMSE appeared in early season (0.15), where sowing and establishment are sensitive to onset timing and early rainfall distribution. Sorghum exhibited a more even profile with a slight mid-season improvement (0.15), consistent with its drought tolerance and broader climatic niche. The “explained variance at peaks” metric ranged from 49% (sorghum) to 62% (rice), indicating that driver emphasis during attention peaks corresponds to substantive reductions in residual uncertainty. These diagnostics corroborate the interpretability claims of the TFT and provide agronomically meaningful narratives to accompany district recommendations.

4.0 CONCLUSION

The results indicated clear and consistent performance gains from multivariate, sequence-based models over classical univariate baselines, with the Temporal Fusion Transformer (TFT) achieving the strongest generalization and interpretability. At an aggregate level, TFT reduced test-fold RMSE to 0.16 t/ha and

increased out-of-sample R^2 to 0.56, outperforming both ARIMA (RMSE $\approx$ 0.26 t/ha; $R^2 \approx$ 0.35) and LSTM (RMSE $\approx$ 0.18-0.21 t/ha; $R^2 \approx$ 0.50). Ablation analyses demonstrated that adding relative humidity (RH) and shortwave radiation to the conventional rainfall-temperature dyad produced further, statistically significant error reductions ($\approx$ 0.02-0.03 t/ha) across crops. District-wise recommendations generated by the TFT engine yielded an average realized gain of about 0.11 t/ha in the test year, with decision-support accuracy near 70%. Together, these findings provide direct evidence that richer climate descriptors and attention-based modeling materially enhance district-level forecasting and operational crop choice.

Interpreted against the research questions, the hierarchy of models suggests that nonlinearity, interactions across covariates, and time-localized effects matter for yield formation in Maharashtra's agro-climates. The negative associations between maximum temperature and yields, coupled with crop-specific salience of radiation and humidity, imply that vapor-pressure dynamics and canopy energy balance are integral to production outcomes rather than secondary details. Attention diagnostics localized importance to phenologically critical windows—establishment for soybean, grain filling for rice, and boll development for cotton—indicating that the timing of weather anomalies, not only their magnitude, governs realized yields. This time-specific sensitivity helps explain why models capable of weighting temporal segments differently (TFT) outperformed those that distribute capacity uniformly across the input window (plain LSTM).

The pattern of gains aligns with established agronomic literature while extending it in two ways. First, traditional studies have emphasized precipitation totals and heat stress; the present results underscore the incremental predictive value of RH and solar radiation in combination with rainfall and temperature. This integration reflects field observations that prolonged humidity deficits or insolation anomalies can disrupt flowering, grain filling, and boll maturation even when seasonal rainfall totals appear adequate. Second, prior modeling work has demonstrated LSTM's advantage over linear models for agricultural time series; here, the additional, consistent edge from transformers suggests that variable-selection networks and multi-head attention provide practically important benefits for multivariate, multi-horizon agri-climate tasks. The district embedding's contribution further supports findings that location-fixed features—soils, irrigation infrastructure, and management routines—shape yield responses and should be represented explicitly rather than implicitly absorbed into error terms.

Each hypothesis is addressed directly by the evidence. H1 is accepted: TFT achieved significantly lower RMSE and higher R^2 than LSTM and ARIMA, with paired tests confirming statistical significance and small-to-medium effect sizes relative to LSTM. H2 is accepted: incorporating RH and radiation improved accuracy over rainfall-temperature baselines across models and crops, with the strongest effects in rice and cotton, consistent with their microclimatic sensitivities. H3 is accepted: district-wise recommendations derived from the best model translated into measurable ex-post yield gains and superior than-chance decision-support accuracy, indicating that forecast improvements were operationally meaningful, not merely statistical refinements.

The broader significance lies in the demonstration that climate-aware crop recommendation can be implemented at statewide scale using routinely available data. By validating improvements under rolling-origin evaluation and perturbation tests, the analysis shows that modest but reliable gains are achievable under realistic data quality, including small gaps and measurement noise. Because the model isolates the marginal contributions of additional drivers, agencies can prioritize data investments strategically. For instance, if humidity and radiation drive gains in specific zones and crops, targeted enhancement of RH sensing or satellite radiation products in those zones may yield outsized returns compared to blanket data augmentation.

Several policy implications follow for Maharashtra's agriculture and allied departments. First, seasonal advisories should integrate RH and radiation alongside rainfall and temperature. Operational bulletins that currently flag rainfall probability and heat alerts can be expanded to include humidity windows and radiation deficits, particularly during flowering and grain filling. Second, district-level recommendation thresholds should be risk-adjusted. The risk-screening rule used here—requiring positive gains at the 25th percentile of predictive distributions—offers a transparent mechanism for extension officers to balance expected benefits against implementation risk. Third, procurement and input distribution can be synchronized with recommendations. Where the advisory indicates a shift towards soybean or cotton in a subset of districts, seed allocation, fertilizer stocking, and credit disbursement may be aligned early to avoid supply frictions that erode realized gains.

Water and insurance policies can also be refined. Irrigation scheduling advisories can be tied to attention-identified windows: for example, supplemental irrigation may be encouraged during mid-season humidity dips that the model flags as yield-critical. Crop-insurance products can incorporate zone-specific model uncertainty, with premium rebates for districts adopting validated recommendations or precision irrigation practices demonstrably correlated with error reductions. In drought-prone interiors, where predictive dispersion remains higher, contingency advisories can specify robust crop choices with lower downside risk even when expected yields are similar across alternatives.

Practical implications for extension and agri-tech providers include the development of explainable dashboards. Variable-importance and attention timelines can be surfaced in simple graphics that show “why” a recommendation is made for a given district and season. This transparency is essential for adoption: when farmers and local officers can see that late-monsoon humidity deficits drive a switch from cotton to sorghum in specific talukas, recommendations become more persuasive and easier to operationalize. Further, integrating district embeddings with local knowledge repositories—soil cards, canal schedules, and historical management—can enhance the credibility and relevance of model outputs without overfitting.

Theoretical implications arise for agri-climate modeling. The consistent incremental value of RH and radiation supports a microclimate-centric view of yield determination that emphasizes boundary-layer processes and canopy energy balance. This suggests that future models should treat these drivers as first-order inputs rather than optional refinements. The superiority of transformer architectures for this task adds to emerging evidence that attention mechanisms map well to agricultural phenology, where the value of information is concentrated in narrow temporal bands. Finally, the empirical utility of static embeddings points to a hybrid theoretical frame where structural, time-invariant heterogeneity interacts with dynamic weather signals to produce realized yields; models that ignore either dimension risk biased attributions and unstable policy guidance.

Notwithstanding these contributions, limitations must be acknowledged. The modeling deliberately excluded explicit market variables, pest incidence, and input-use intensity, relying instead on cropping-pattern shares and district embeddings as partial proxies. While this choice allowed clean identification of the incremental climate signal, it leaves room for omitted-variable effects where management practices vary abruptly. Annual yield targets also compress intra-seasonal variability; although sequence inputs captured monthly dynamics, end-of-season shocks and post-harvest losses may not be fully reflected. Satellite vegetation indices, which often improve early-season prediction, were excluded to isolate the contribution of humidity and radiation; their integration would likely raise accuracy further but could confound attribution if not introduced in controlled ablations. Finally, recommendation gains were validated in a single recent test year; multi-year operational pilots would strengthen external validity and reveal adoption frictions.

Future research can proceed along five avenues. First, integrating vegetation indices and microwave soil-moisture products in a staged manner can test whether biophysical state variables complement or substitute for RH and radiation in different zones. Second, multi-objective optimization that balances expected yield with downside risk and water demand could refine recommendations in water-stressed districts. Third, causal designs—such as difference-in-differences using phased advisory rollouts—could estimate realized impact under adoption, disentangling model quality from implementation constraints. Fourth, ensemble approaches that blend TFT with gradient-boosted trees for static effects, or with drought indices for interiors, may reduce dispersion where volatility is largest. Fifth, human-centered evaluations with extension personnel can assess interpretability assets, comparing attention-based explanations against simpler rule lists to identify formats that best drive behavioral change.

In summary, the analysis has shown that attention-based, multivariate time-series models using an enriched climate feature set can deliver statistically and operationally meaningful gains in district-level yield prediction and crop recommendation for Maharashtra. The results indicate that policy and practice should update advisory content to include humidity and radiation, adopt risk-adjusted thresholds for implementation, and deploy explainable tools that connect climate drivers to phenological windows. While further work is needed to incorporate market and biotic stressors and to evaluate adoption at scale, the present evidence is sufficient to justify a phased operational deployment with monitoring, enabling responsive, climate-smart crop planning in a state where weather variability is both a persistent risk and an actionable signal.

5.0 REFERENCES

1. Bhat, R., Sreeganga, S. D., & Priyanka, P. (2022). Crop production prediction models in Indian agriculture: A review. *Indian Ecological Society Review*, 2(1), 1–13.
2. Chipade, A. M., & Salunkhe, S. (2024). Impact of climate variability on crop yield in Pune district. *E3S Web of Conferences*, 550, 01012.
3. Deforce, B. (2024). Harnessing the power of transformers and data fusion in smart agriculture. *Applied Soft Computing*, 153, 110084.
4. Deforce, B., Van Herck, J., Baert, A., & De Moor, B. (2022). Forecasting sensor-data in smart agriculture with Temporal Fusion Transformers. *KU Leuven Working Paper*.
5. IMD. (2024/2025). IMD new high spatial resolution ($0.25^\circ \times 0.25^\circ$) long period (1901–2024) daily gridded rainfall data set over India (product page). India Meteorological Department.
6. Javed, M. A., Paul, M., & Rahman, M. M. (2024). Crop yield prediction in agriculture: A systematic literature review. *Heliyon*, 10(11), e3168673.
7. Jena, P., Chattopadhyay, R., Kundu, P., Maiti, P., Rajeevan, M., & Pai, D. S. (2020). Performance analysis of IMD high-resolution gridded rainfall data over India. *Journal of Hydrometeorology*, 21(7), 1761–1775.
8. Junankar, V., et al. (2023). Wheat yield prediction using Temporal Fusion Transformers. Preprint/Conference record.
9. Kalmani, H. V., Dharwadkar, V. N., & Thapa, V. (2025). Crop yield prediction using deep learning algorithm based on CNN-LSTM with attention layer and skip connection. *Indian Journal of Agricultural Research*, 59(8), 1303–1311.
10. Kelkar, S. M., Tambe, S. S., & Kulkarni, P. (2020). Impact of climate variability and change on crop production in Maharashtra, India. *Current Science*, 118(8), 1235–1243.
11. Khedekar, S. Y., Panchabhai, S. N., & Mandal, T. K. (2023). Study on effect of rainfall variability on crop productivity of tur (arhar) and cotton crops for various districts of Maharashtra, India. *International Journal of Environment and Climate Change*, 13(9), 2347–2357.
12. Kumar, M., & Rana, S. K. (2025). A comparative statistical approach for forecasting maize yield in India using ARIMA models. *Asian Journal of Probability and Statistics*, 27(9), 1–9.
13. Lim, B., Arik, S. O., Loeff, N., & Pfister, T. (2019/2021). Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *Nature Machine Intelligence*.
14. Maldonado-Cruz, E., Ortega-Torres, J., Tavera-Mendoza, L., & Esquivel-Gómez, J. (2024). Multi-horizon well performance forecasting with Temporal Fusion Transformer. *Energy Reports*, 10, 100339.
15. Pai, D. S., Sridhar, L., Rajeevan, M., Sreejith, O. P., Satbhai, N. S., & Mukhopadhyay, B. (2014). Development of a new high spatial resolution ($0.25^\circ \times 0.25^\circ$) daily gridded rainfall data set over India—Analysis of break and active spells of the southwest monsoon. *MAUSAM*, 65(1), 1–18.
16. Peringiyil, A., Raval, L., Mishra, V., & Jain, S. (2025). A station-based 0.1-degree daily gridded ensemble precipitation dataset for India. *Scientific Data*, 12, Article 123.
17. Prasad, G., Sinha, S., & Singh, G. (2018). Impact of temperature and solar radiation on wheat crop productivity. *Vayu Mandal*, 44(2), 7–17.
18. Saksena, R., Kanani, H., & Patidar, S. (2023). Optimizing LSTM and Bi-LSTM models for crop yield prediction. *Applied Intelligence*, 53(18), 22179–22196.
19. Shawon, S. M. A. K., Sumi, S. M. A., Naser, F., & Rahman, M. M. (2024). Crop yield prediction using machine learning: An extensive review. *Computational Agriculture*, 2, 100061.
20. Srivastava, A. K., Rajeevan, M., & Kshirsagar, S. R. (2009). Development of a high-resolution daily gridded temperature data set (1969–2005) for the Indian region. *Met monograph/IMD*.
21. Subramaniam, L. K., Chithambarathanu, T., & Periasamy, R. (2024). Crop yield prediction using effective deep learning and dimensionality reduction approaches for Indian regional crops. *Agricultural Informatics*, 5(3), 100191.
22. Times of India. (2025a, October 31). Scientists warn of 50% drop in grape yield across Maharashtra this season. *The Times of India*.
23. Times of India. (2025b, August 1). Maha banks have over ₹35,000 cr farm loan non-performing assets. *The Times of India*.
24. Xie, W., Zhao, B., Wang, J., & Li, Y. (2024). Recent advances in transformer technology for agriculture: A comprehensive survey. *Engineering Applications of Artificial Intelligence*, 136, 108980.
25. Zhong, L., Hu, L., & Zhou, H. (2020). Deep learning based multi-temporal crop classification and yield prediction using CNN-RNN. *Computers and Electronics in Agriculture*, 177, 105704002E