Natural Resources for Human Health
eISSN: 2583-1194

Original Research

Received 13 July 2026
Revised 23 Aug 2026
Accepted 11 Sep 2026

Natural Resources for Human Health 2026, Vol. 6 (20s): 1043–1067

KEYWORDS:
soil moisture forecasting, machine learning, persistence baseline, smart agriculture, agricultural sensing, time-series forecasting, temporal validation, block bootstrap, robustness analysis.

Beyond Persistence: A Chronological Benchmark of Machine Learning for Short-Term Soil Moisture Forecasting Across Station-Year Datasets

Daggumati Pavan Hari Krishna[1], Sonti Thirupathamma[2], M. V. D. Prasad[1]

[1] Department of Computer Science and Engineering, K L Deemed to be University, Vaddeswaram, Andhra Pradesh, India.
[2] Department of Electronics and Communication Engineering, SRM University–AP, Amaravati, Andhra Pradesh, India. Ph.D. Scholar.
Corresponding Author: Sonti Thirupathamma
Ph.D. Scholar, Department of Electronics and Communication Engineering, SRM University–AP, Amaravati, Andhra Pradesh, India.
Email: thirupathamma_sonti@srmap.edu.in

ABSTRACT: Short-term soil moisture forecasting can support agricultural monitoring and irrigation decisionsupport systems, but the incremental value of machine learning (ML) should be evaluated against strong temporal baselines. This study asks whether ML provides predictive value beyond persistence for shortterm soil-moisture forecasting across station-year datasets and multiple horizons. A publicly available AgriDataValue environmental dataset containing soil moisture, weather, and soil-temperature observations was evaluated across two monitoring stations and four station-year datasets. The raw dataset contained 22,111 observations; quality control retained 22,050 observations (99.724%). Ridge regression, Random Forest, and histogram-based gradient boosting (HGB) were compared with a persistence forecast at 1-, 6-, 12-, and 24-hour horizons. Samples were constructed only within contiguous temporal sequences, and targets were retained only when the future timestamp occurred at the exact requested elapsed time. Chronological expanding-window validation used an initial 50% training window followed by two sequential test folds. MAE, RMSE, and R^2 were reported, while temporal block bootstrap resampling with 2,000 replicates quantified uncertainty in ML-minus-persistence absolute-error differences. Persistence achieved the lowest MAE in all eight station-horizon combinations. Ridge was the strongest evaluated ML model but remained worse than persistence at every horizon. The mean relative MAE increase of the best ML model across the eight station-horizon combinations was 117.68%, with a median increase of 103.59%. All primary Ridge-minus-persistence 95% bootstrap confidence intervals were strictly above zero. Robustness analyses included seasonal persistence, a controlled LSTM baseline, grouped feature ablation, training-window sensitivity, bootstrap block-length sensitivity, and sequence-level ADF/KPSS stationarity analysis. These analyses did not alter the primary conclusion. Under the evaluated dataset, feature configuration, model families, horizons, and chronological validation protocol, the tested ML approaches did not demonstrate incremental predictive value beyond persistence. The result is specific to the investigated setting and is not evidence that ML is universally ineffective for soil-moisture forecasting.

Natural Resources for Human Health
eISSN: 2583-1194

1. INTRODUCTION

Soil moisture is a fundamental environmental state variable for agricultural monitoring because it reflects water availability in the soil and is related to crop growth, water stress, and irrigation management. The deployment of environmental sensors has enabled high-frequency observations of soil moisture together with meteorological variables such as air temperature, humidity, pressure, precipitation, and dew point. These measurements create opportunities for data-driven forecasting systems that can contribute to agricultural decision support.

Machine learning has become an important methodological direction for soil-moisture prediction. Recent reviews show extensive use of regression, ensemble learning, recurrent neural networks, hybrid architectures, and remote-sensing-based methods, but also substantial variation in datasets, sensor technologies, input variables, forecasting horizons, and validation strategies [1], [2], [8]. Field studies have reported useful performance for boosting, Random Forest, XGBoost, and deep-learning models under particular agricultural conditions [3], [4], [6], [7]. These results demonstrate the potential of ML, but they do not establish that ML necessarily adds predictive information beyond a strong temporal baseline in every dataset.

This issue is particularly important for soil moisture because the target can be strongly temporally dependent. At short horizons, the latest observed value may already contain much of the information required for a useful forecast. A persistence forecast formalizes this behavior by assigning the current observation to the future target. If an ML model cannot improve upon persistence, its additional features, nonlinear representation, and training complexity do not establish incremental predictive value for that forecasting task.

The central research question (RQ) is:

Does machine learning provide incremental predictive value over a persistence baseline for short-term soil-moisture forecasting across different forecasting horizons and agricultural monitoring environments?

The secondary questions are: (i) how does the relative ML–persistence difference change with forecasting horizon; (ii) is the observed difference robust to temporal dependence and uncertainty; and (iii) do alternative baselines, model families, feature groups, training-window choices, bootstrap block lengths, and stationarity diagnostics alter the primary conclusion?

The contributions are:

1) persistence-aware evaluation of ML rather than ML accuracy in isolation;

2) four forecasting horizons across two stations and four station-year experimental units;

3) chronological expanding-window validation to avoid future-to-past leakage;

4) temporal block bootstrap uncertainty analysis with 2,000 resamples;

5) a broad robustness suite covering seasonal persistence, a controlled LSTM, feature ablation, training-window sensitivity, bootstrap block-length sensitivity, and stationarity diagnostics; and

6) transparent reporting of a negative empirical result without generalizing it beyond the investigated experimental conditions.

2. RELATED WORK

2.1 Sensor-Based Soil Moisture Forecasting

Ivanova reviewed 60 scientific articles on sensor-based soil-moisture forecasting published from 2014 to 2024 and highlighted the lack of a transparent, unified methodology for comparing forecasting systems [1]. Sensor-based systems commonly combine recent soil moisture with weather variables and historical observations. Their performance is affected by measurement quality, temporal continuity, sensor placement, environmental variability, and forecast horizon.

2.2 Machine Learning and Ensemble Models

Togneri *et al.* evaluated soil-moisture forecasting across twelve fields, four farms, eight crop types, and four years and reported strong performance for boosting-based approaches [3]. Custodio´ and Prati compared modern and traditional methods and

reported cases in which traditional univariate methods outperformed modern ML/deep-learning approaches [4]. The latter result is especially relevant to the present study because it illustrates why model complexity should be assessed relative to strong temporal baselines.

2.3 Deep Learning and Hybrid Models

Recent work has investigated GRU–Transformer architectures, LSTM models, and hybrid approaches for soil-moisture prediction [5]–[7]. Su *et al.* reported strong daily cotton forecasting with a minimal-input LSTM and transfer learning across growing seasons [7]. Tefera *et al.* combined ground sensors with biometeorological variables and evaluated RF, XGBoost, neural networks, and LSTM models in Texas agroecosystems [6]. These studies indicate that sequenceaware and feature-rich models can be valuable in suitable environments; they also reinforce the importance of reporting the baseline against which that value is measured.

2.4 Persistence and Traditional Baselines

Persistence is a deliberately simple forecast that assumes the current state remains unchanged over the forecast horizon. For slowly varying environmental variables, persistence can be competitive at short horizons. Therefore, a model should demonstrate improvement over persistence rather than only report a high R^2 or low absolute error.

2.5 Research Gap

The specific gap addressed here is the systematic measurement of incremental ML value relative to persistence under one common chronological, multi-horizon framework with dependence-aware uncertainty estimation. The study therefore focuses on the question of whether added model complexity produces measurable forecasting improvement, rather than assuming that it does.

3. DATASET AND DATA PREPARATION

3.1 Dataset Source and Experimental Units

The study uses the public "AgriDataValue – Environmental Data Useful for Smart Irrigation" dataset archived on Zenodo (DOI: 10.5281/zenodo.18959383) [9]. The dataset contains environmental observations including air temperature, air humidity, air pressure, dew point, precipitation, earth humidity (soil moisture), and earth temperature.

Four station-year datasets were analyzed:

7) Station 34, 2024 – Sorghum (Jumbo Star);

8) Station 35, 2024 – Corn (Forturio); 3) Station 34, 2025 – Corn; and 4) Station 35, 2025 – Corn.

The raw total was 22,111 observations. Quality control removed 61 observations, leaving 22,050 observations (99.724%). The station-year dataset is the experimental unit; individual observations are forecast samples within each unit.

Table 1. Dataset quality-control summary.

Dataset	Crop	Raw	Removed	Clean	Retained	Sequences
S34–2024	Sorghum	6738	1	6737	99.985%	6
S35–2024	Corn	4809	7	4802	99.854%	10
S34–2025	Corn	6378	5	6373	99.922%	18
S35–2025	Corn	4186	48	4138	98.853%	27
Total	–	22111	61	22050	99.724%	–

3.2 Quality Control

Quality control was performed before feature construction. Duplicate timestamps were removed. Soil-moisture values outside the recorded range were checked, and physically suspicious environmental measurements were examined. Large one-hour soil-moisture changes were treated as diagnostic indicators rather than automatically deleting every large change. Observations lacking required model inputs were excluded from the corresponding forecasting sample. No imputation was used to bridge temporal sequence breaks, and no future target was created across a large missing-data gap.

Exact 100% soil-moisture values were retained because they are recorded observations and are not automatically invalid. The largest number of exclusions occurred in Station 35 during 2025, where 48 observations were removed from 4,186 raw records.

The same quality-control principles were applied across all four station-year datasets to maintain a consistent basis for comparison. Quality control was designed to remove invalid records without suppressing genuine environmental variability; consequently, large changes were inspected diagnostically rather than discarded solely because of their magnitude.

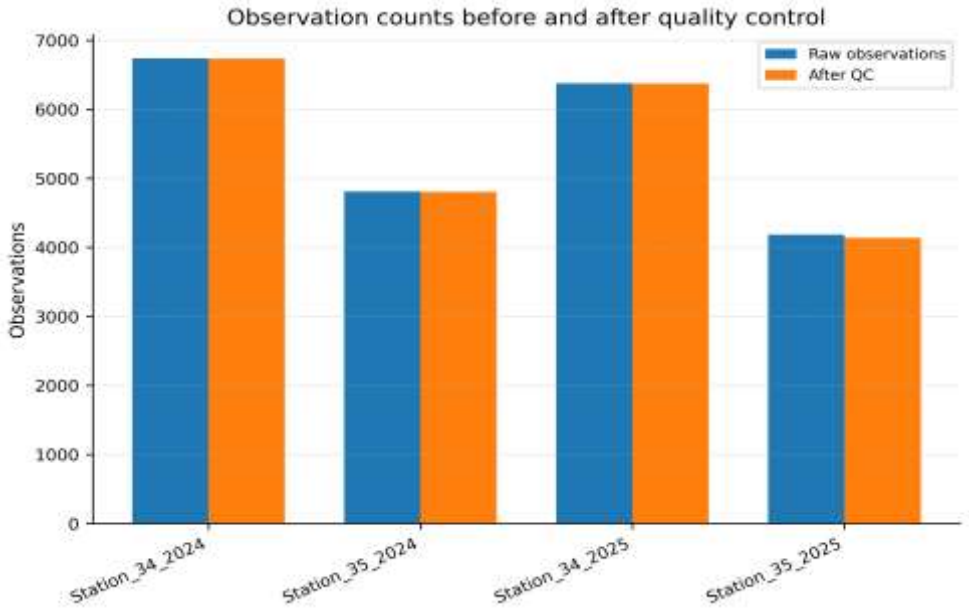


Fig. 1: Raw observations, quality-control exclusions, and retained observations for all four station-year datasets.

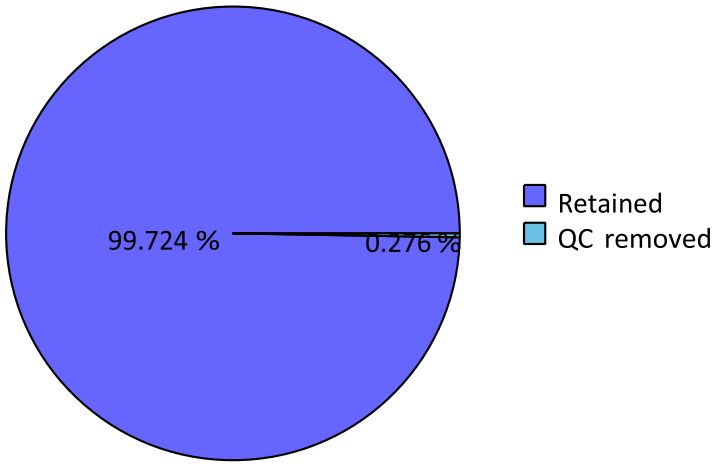


Fig. 2: Overall quality-control retention composition: 22,050 retained observations and 61 excluded observations from 22,111 raw observations.

Natural Resources for Human Health
eISSN: 2583-1194

3.3 Temporal Coverage and Distribution Diagnostics

Station 34–2024 contained 6,737 cleaned observations from 2024-03-07 10:00 to 2024-12-31 00:00, with 99.90% exact one-hour intervals, six gaps greater than one hour, and a largest gap of 392 h. Station 35–2024 contained 4,802 observations from 2024-03-15 12:00 to 2024-12-18

16:00, with 99.92% exact one-hour intervals, four gaps, and a largest gap of 1,818 h. Station 34–2025 contained 6,373 observations from 2025-01-01 01:00 to 2025-10-13 00:00, with 99.76% exact one-hour intervals, fifteen gaps, and a largest gap of 424 h. Station 35–2025 contained 4,138 observations from 2025-01-08 11:00 to 2025-10-13 00:00, with 99.86% exact one-hour intervals, six gaps, and a largest gap of 1,763 h.

3.4 Temporal Sequence Construction

The observations were divided into contiguous temporal sequences. Forecasting samples were generated independently within each sequence. For horizon h, a target was accepted only when the future timestamp existed exactly h hours after the current timestamp:

$$y_{t+h} = \text{soil moisture at } (t + h). \quad (1)$$

This prevented artificial target creation across long gaps and ensured that a 6-, 12-, or 24-hour target represented the requested elapsed time rather than an arbitrary next available observation.

3.5 Feature Engineering

The primary ML input contained 40 features: current soil moisture; soil-moisture lags at 1, 3, 6, 12, and 24 h; 1- and 6-h soil-moisture changes; current air temperature, humidity, pressure, dew point, precipitation, and soil temperature; 1-, 6-, 12-, and 24-h lags of air temperature, humidity, dew point, precipitation, and soil temperature; 6- and 24-h precipitation sums; and sine/cosine encodings of hour-of-day and day-of-year.

The complete feature vector was: The 40-feature representation consisted of five groups: (i) current soil moisture, soil-moisture lags at 1, 3, 6, 12, and 24 h, and soil-moisture change features; (ii) current air temperature, humidity, pressure, dew point, precipitation, and soil temperature; (iii)

1-, 6-, 12-, and 24-h lags of the principal meteorological variables; (iv) 6- and 24-h precipitation

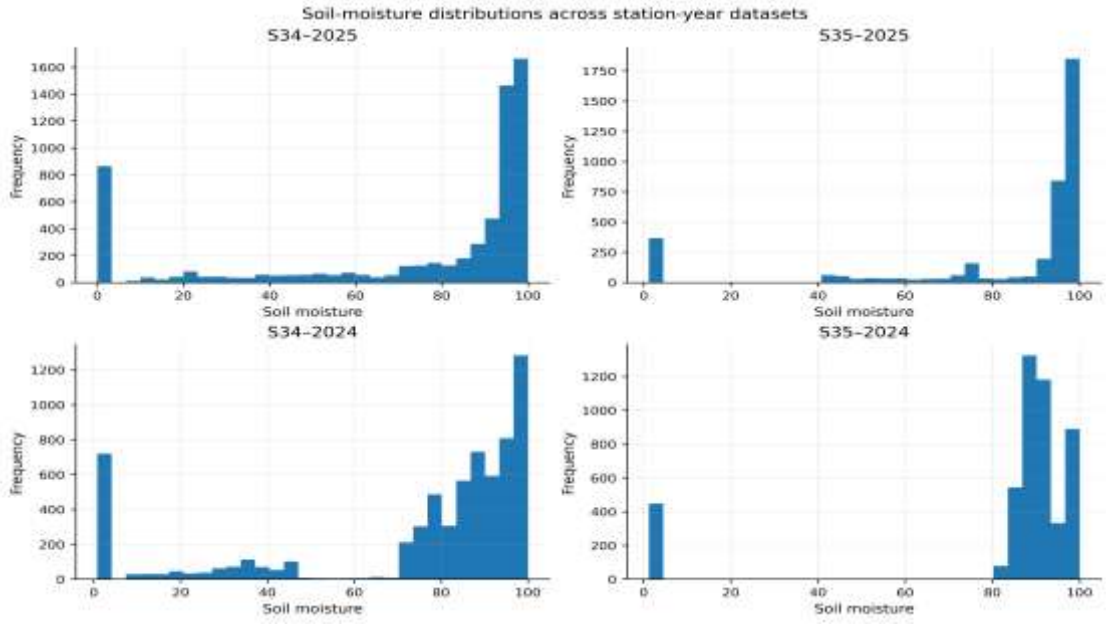


Fig. 3: Soil-moisture distributions for the four station-year datasets. The distributions illustrate the observed range and concentration of the target variable after quality control

Natural Resources for Human Health
eISSN: 2583-1194

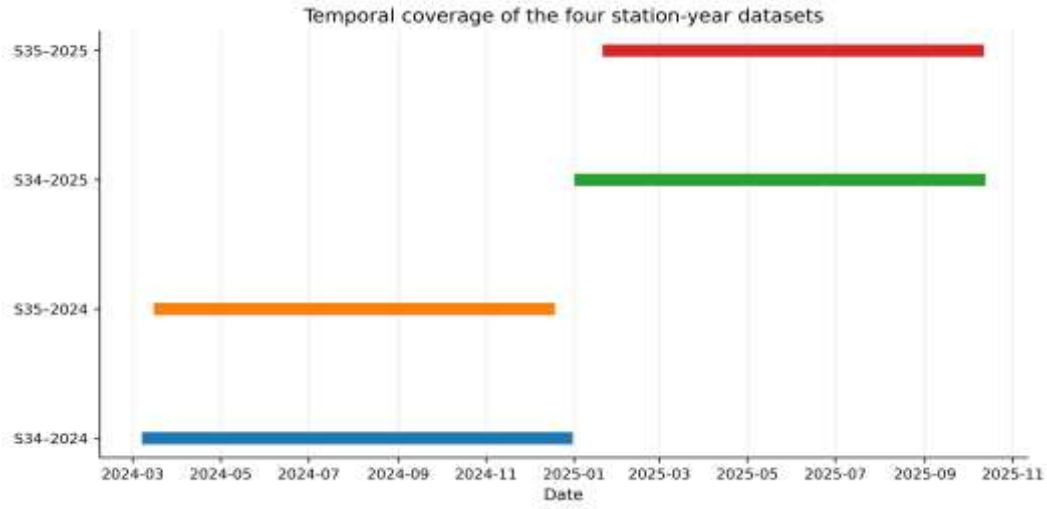


Fig. 4: Temporal coverage of the four station-year datasets. Large measurement gaps were retained as sequence boundaries rather than bridged by interpolation.

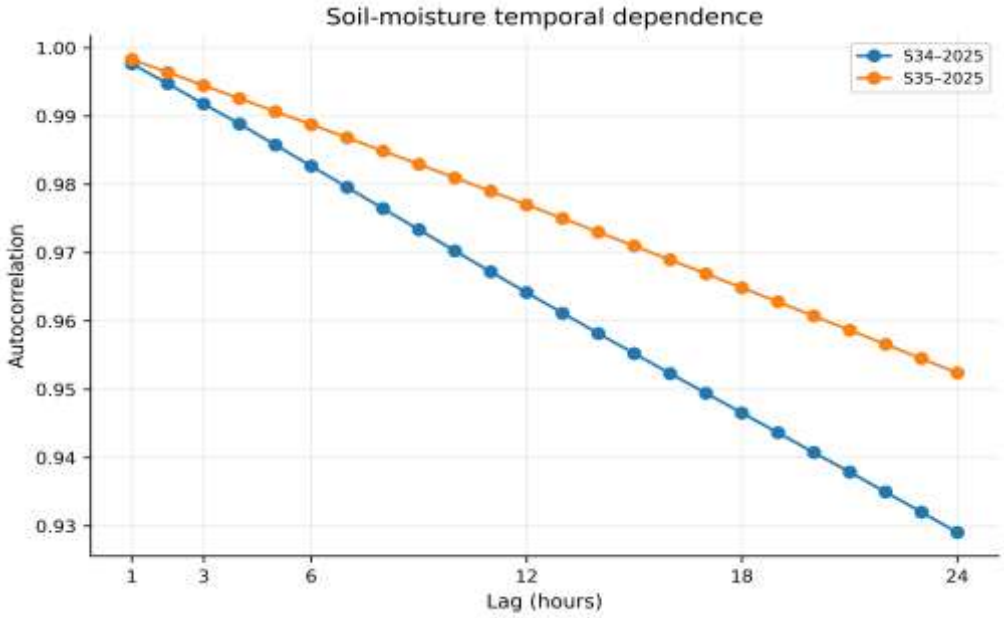


Fig. 5: Soil-moisture autocorrelation diagnostics. Strong short-lag dependence motivates the use of persistence as a primary forecasting baseline.

accumulation features; and (v) cyclical hour-of-day and day-of-year encodings. This representation combines recent soil state, meteorological context, temporal history, and seasonal timing without using future observations.

Cyclical variables were represented as

$$x_{\sin} = \sin\left(\frac{2\pi x}{P}\right), \qquad x_{\cos} = \cos\left(\frac{2\pi x}{P}\right), \tag{2}$$

where P is the relevant period.

Natural Resources for Human Health
eISSN: 2583-1194

4. EXPERIMENTAL DESIGN

4.1 Forecasting Horizons

Four horizons were evaluated: 1, 6, 12, and 24 h. These horizons were selected to represent near-term and progressively longer short-term forecasts while retaining sufficient observations for chronological testing.

4.2 Persistence Baseline

For a horizon h, persistence was defined as

$$\hat{y}^{P}_{i+h} = y_{t,} \tag{3}$$

where y_t is the latest observed soil-moisture value. Persistence requires no model fitting or parameter estimation.

4.3 Primary Machine-Learning Models and Hyperparameters

Ridge regression was used as a regularized linear model with $a = 10$. Random Forest used 200 trees, maximum depth 12, minimum leaf size 5, random state 42, and parallel training. HGB used 300 boosting iterations, learning rate 0.05, maximum 31 leaf nodes, L_2 regularization 1, and random state 42. The implementation used the scikit-learn ecosystem.

Table 2. Primary model configurations.

Model	Configuration
Ridge	$a = 10$
Random Forest	200 trees; max depth 12; min samples leaf 5; random state 42
HGB	300 iterations; learning rate 0.05; max leaf nodes 31; $L_2 = 1$; random state 42

4.4 Chronological Expanding-Window Validation

Random train-test splitting was avoided because neighboring observations are temporally dependent. Approximately the first 50% of available forecasting samples formed the initial training window. The remaining 50% were divided into two sequential test folds. After evaluation of the first test fold, that fold was added to the training set before evaluating the second fold. No future test observations were used to train an earlier forecast.

This ordering preserves the forecasting direction of time: each test fold is evaluated using only observations available before that fold, and the earlier test fold is incorporated only before the later fold is evaluated. Accordingly, the reported estimates represent chronological forecasting performance rather than performance under a randomly shuffled split.

4.5 Evaluation Metrics

The primary metric was mean absolute error (MAE):

$$MAE = \frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i| \tag{4}$$

RMSE was calculated as

$$RMSE = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \tag{5}$$

The coefficient of determination was

$$R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y})^2}. \quad (6)$$

Relative skill against persistence was defined as

$$Skill = 1 - \frac{MAE_{ML}}{MAE_P}, \quad (7)$$

where positive values indicate improvement over persistence and negative values indicate worse performance.

The relative MAE increase was

$$Increase(\%) = 100 \times \frac{MAE_{ML} - MAE_P}{MAE_P}. \quad (8)$$

4.6 Temporal Block Bootstrap

For each ML model and horizon, the observation-level absolute-error difference was defined as

$$\Delta e_i = |y_i - \hat{y}_i^{ML}| - |y_i - \hat{y}_i^{P}|. \quad (9)$$

Positive Δe_i means that the ML model made a larger absolute error than persistence. The primary inference used 24-h circular temporal blocks and 2,000 bootstrap replicates. The 2.5th and 97.5th percentiles defined the 95% confidence interval. Sensitivity was additionally assessed using 12-, 24-, 48-, and 168-h blocks [10].

The use of contiguous circular blocks retains short-range temporal dependence within resampled error sequences, unlike an observation-wise bootstrap that would treat neighboring errors as independent. A confidence interval entirely above zero indicates that the ML model's absolute error exceeded persistence's error over the evaluated comparison under the specified bootstrap procedure; it should be interpreted as evidence for this experimental setting rather than as a universal statement about other datasets.

4.7 Stationarity Diagnostics

Sequence-level Augmented Dickey–Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests were applied. The tests have different null hypotheses and were therefore interpreted jointly rather than treating either test as a definitive classification [11], [12].

Specifically, the ADF test evaluates a unit-root null hypothesis, whereas KPSS evaluates a stationarity null hypothesis. Interpreting both tests together helps distinguish consistent evidence from conflicting or inconclusive sequence-level results; the tests were not used to label an entire station-year from a single statistic.

4.8 Robustness Design

The robustness suite tested: (1) seasonal persistence at 24, 48, and 168 h; (2) a controlled

LSTM with a 24-h input sequence and 11 physical/time features; (3) grouped feature ablations;

(4) Ridge 24-h sensitivity to 40%, 50%, and 60% initial training windows; (5) 12-, 24-, 48-, and

168-h bootstrap block lengths; and (6) ADF/KPSS stationarity across all four station-year datasets. These analyses were treated as robustness checks rather than opportunities for post-hoc tuning.

5. PRIMARY RESULTS

5.1 MAE

Persistence achieved the lowest MAE in all eight station-horizon combinations. Ridge was the strongest ML model by MAE in all eight comparisons.

Table 3. Primary MAE results.

Station	h	Persistence	Ridge	RF	HGB
34	1	0.2737	0.4982	7.2491	10.7513
34	6	1.4159	2.4982	11.0179	16.2134
34	12	2.1543	4.4861	13.8255	17.3368
34	24	2.4435	8.3801	12.4970	18.6861
35	1	0.1293	0.2541	12.3218	12.5087
35	6	0.6231	1.2475	12.6895	13.8525
35	12	1.1197	2.3177	13.3439	13.7994
35	24	1.9945	4.5491	13.6477	13.8036

5.2 RMSE

Persistence also produced the lowest RMSE in all eight comparisons.

5.3 R^2

Persistence retained high R^2 values at all horizons. At 24 h, persistence achieved $R^2 = 0.906937$ at Station 34 and 0.926658 at Station 35. Ridge achieved 0.749894 and 0.887081, respectively. HGB produced negative R^2 at 12 and 24 h for Station 34.

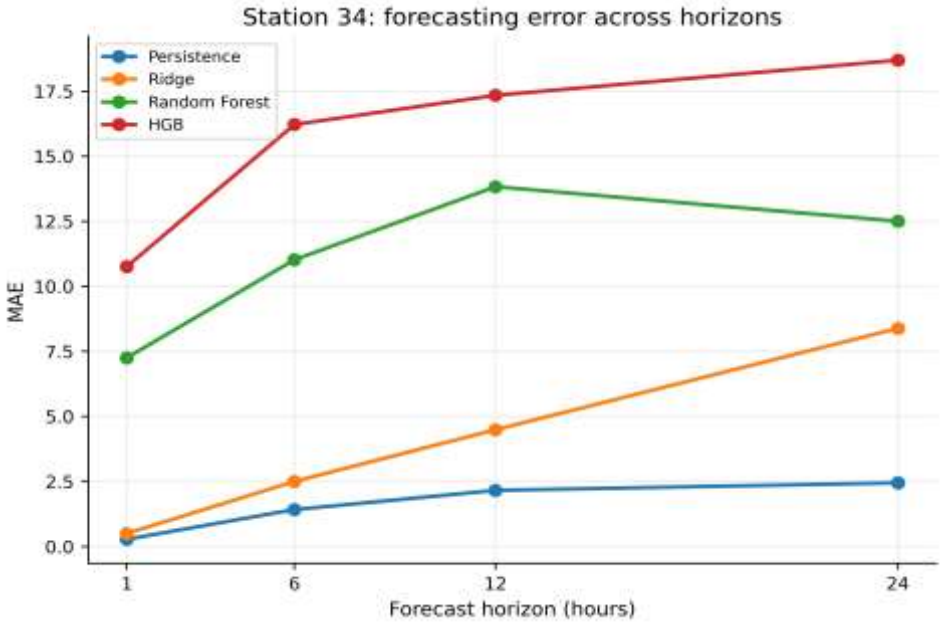


Fig. 6: Station 34 MAE across 1-, 6-, 12-, and 24-h horizons. Persistence remains below every evaluated ML model.

Table 4. Primary RMSE results.

Station	h	Persistence	Ridge	RF	HGB
34	1	1.0018	1.1749	11.9199	17.3184
34	6	3.7472	4.4665	16.5622	22.3329

34	12	5.4230	7.0038	19.9051	24.4632
34	24	6.6711	10.4621	17.9906	25.3449
35	1	0.9361	0.9917	20.0930	20.2181
35	6	2.9644	3.3533	20.4551	21.0966
35	12	4.3601	5.2331	21.2403	21.7084
35	24	6.4969	8.7321	21.5039	21.6176

Table 5. Primary R2 results.

Station	h	Persistence	Ridge	RF	HGB
34	1	0.9981	0.9976	0.7095	0.4680
34	6	0.9729	0.9646	0.4013	0.1036
34	12	0.9418	0.9075	0.1567	-0.1680
34	24	0.9069	0.7499	0.1729	-0.3761
35	1	0.9985	0.9983	0.5177	0.5146
35	6	0.9854	0.9825	0.5039	0.4878
35	12	0.9681	0.9589	0.4783	0.4638
35	24	0.9267	0.8871	0.4632	0.4547

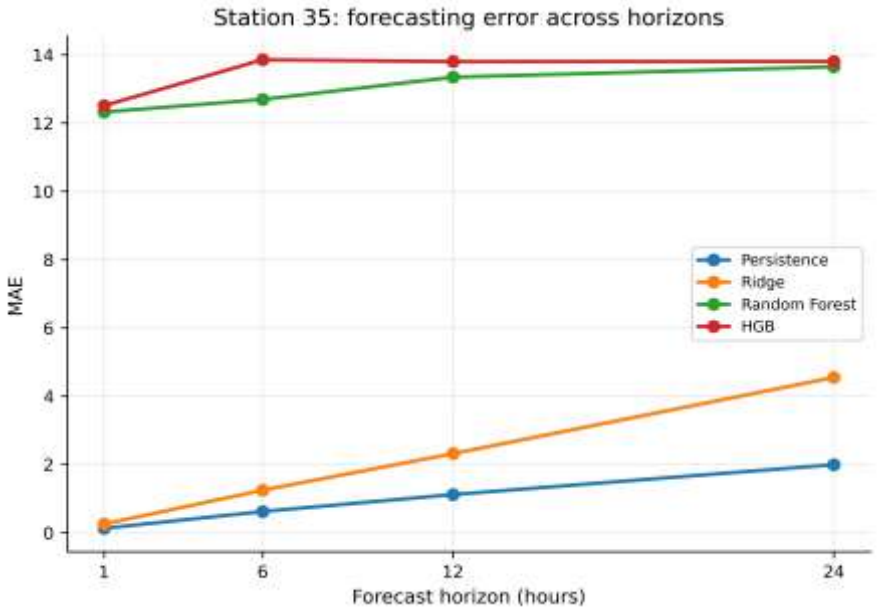


Fig. 7: Station 35 MAE across 1-, 6-, 12-, and 24-h horizons. Persistence remains below every evaluated ML model.

Natural Resources for Human Health
eISSN: 2583-1194

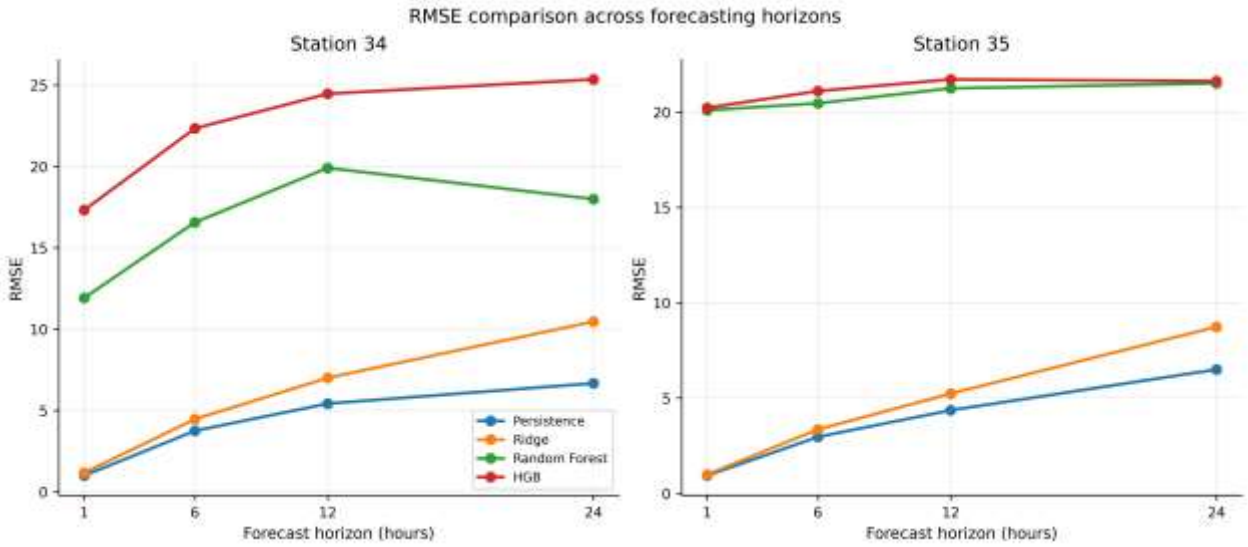


Fig. 8: RMSE for persistence, Ridge, Random Forest, and HGB across both stations and all four horizons.

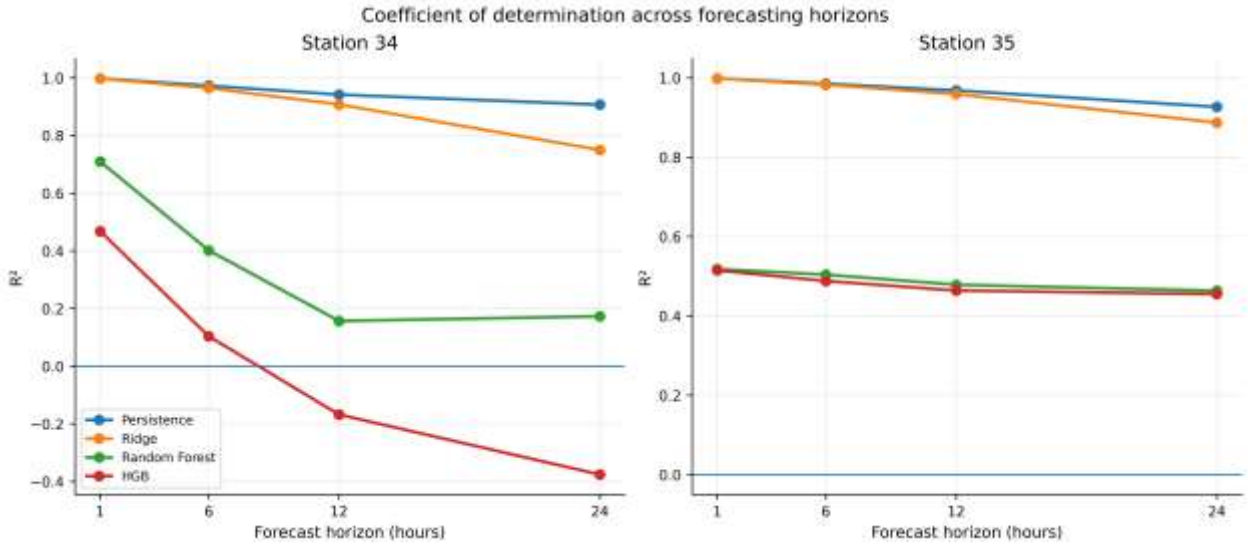


Fig. 9: Coefficient of determination across models, stations, and forecasting horizons.

5.4 Best ML Model Versus Persistence

Ridge was the best ML model in all eight station-horizon combinations. The mean relative MAE increase was 117.68%, while the median increase was 103.59%.

Table 6. Ridge versus persistence with bootstrap uncertainty.

St.	h	MAE_P	MAE_R	ΔMAE	95% CI	Increase
34	1	0.2737	0.4982	0.2245	[0.1704,0.2775]	82.00%
34	6	1.4159	2.4982	1.0823	[0.8938,1.2543]	76.44%
34	12	2.1543	4.4861	2.3319	[1.9480,2.6724]	108.24%
34	24	2.4435	8.3801	5.9366	[5.1566,6.6952]	242.95%

35	1	0.1293	0.2541	0.1248	[0.0963,0.1544]	96.50%
35	6	0.6231	1.2475	0.6244	[0.4481,0.8097]	100.20%
35	12	1.1197	2.3177	1.1984	[0.8657,1.5636]	106.99%
35	24	1.9945	4.5491	2.5546	[1.9148,3.2902]	128.08%

5.5 Bootstrap Inference Across All ML Models

All 24 primary ML-minus-persistence bootstrap comparisons across three ML models, two stations, and four horizons produced 95% confidence intervals entirely above zero. Thus, the primary result was not limited to Ridge.

6. ROBUSTNESS AND DIAGNOSTIC RESULTS

6.1 Seasonal Persistence

Seasonal persistence at 24, 48, and 168 h was consistently worse than ordinary persistence in the 2025 robustness datasets. At Station 34 and the 24-h forecast, ordinary persistence MAE was 3.6720 versus 7.1660, 10.7151, and 21.3185 for 24-, 48-, and 168-h seasonal persistence. At Station 35, the corresponding values were 1.9191 versus 3.8600, 5.7509, and 12.2985.

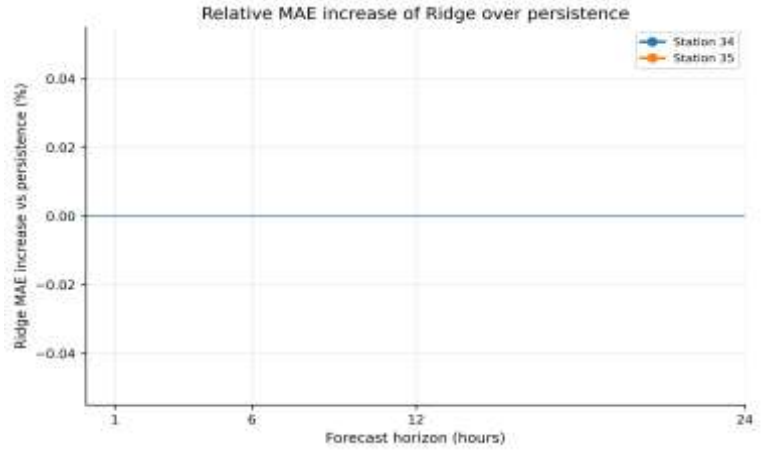


Fig. 10: Relative MAE increase of Ridge over persistence. Positive values indicate higher Ridge error.

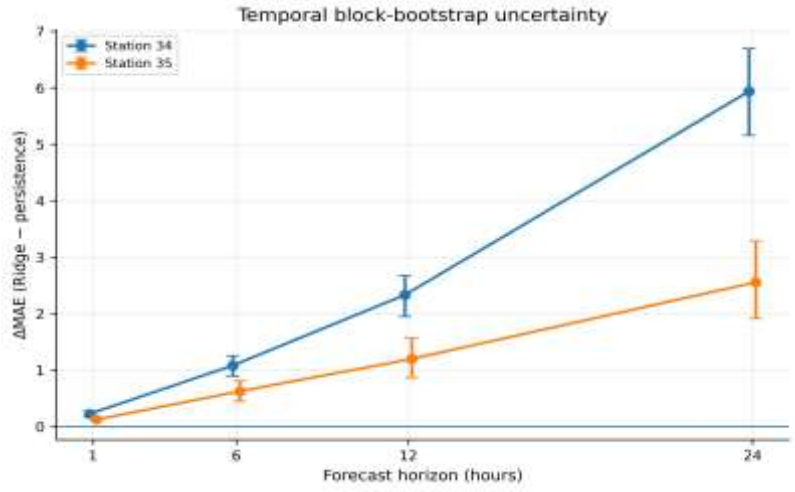


Fig. 11: Temporal block-bootstrap confidence intervals for ML-minus-persistence absolute-error differences. Positive values indicate higher ML error.

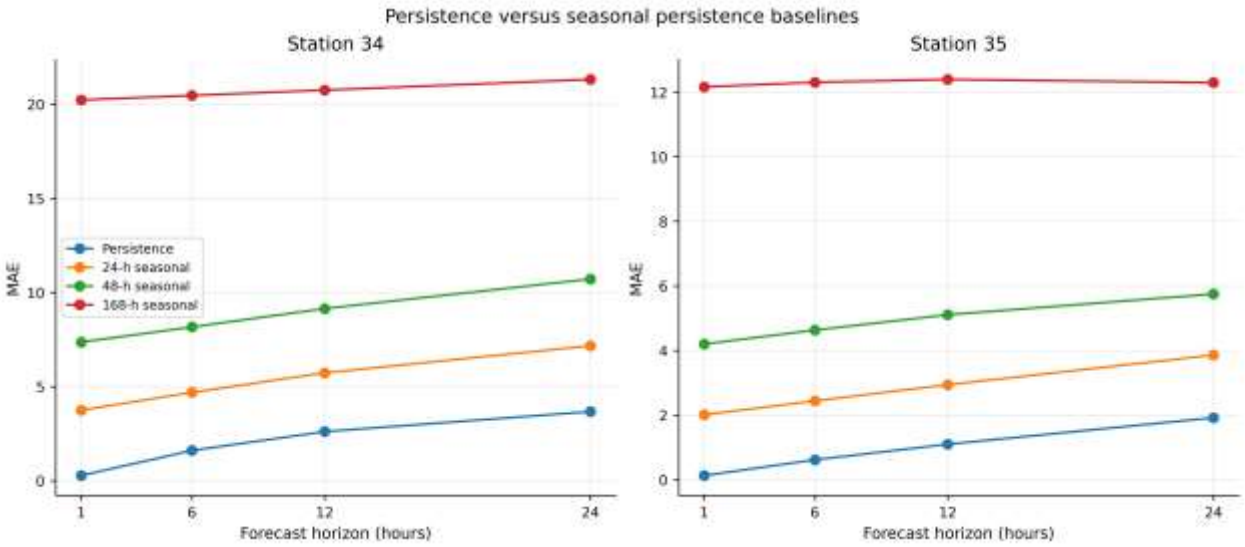


Fig. 12: Seasonal persistence MAE compared with ordinary persistence at 24, 48, and 168 h.

6.2 Controlled LSTM Baseline

The controlled LSTM used 24-h input sequences, 11 physical/time features, a 32-unit LSTM layer, dropout 0.2, a 16-unit dense layer, standardized inputs, a MinMax-scaled target, sigmoid output, Adam learning rate 0.0005, gradient clipping at 1, batch size 64, up to 25 epochs, and early stopping patience 4. The LSTM did not outperform persistence at any evaluated horizon.

Table 7. Controlled LSTM robustness results.

St.	h	LSTM MAE	Persistence MAE	Skill (%)
34	1	20.4781	0.2672	-175.57
34	6	18.6061	1.3801	-25.87
34	12	18.9974	2.2054	-12.51
34	24	32.5572	2.6539	-13.55
35	1	28.2969	0.1293	-276.06
35	6	27.1401	0.6233	-53.20
35	12	26.5679	1.1191	-26.71
35	24	29.4180	1.9934	-16.65

6.3 Grouped Feature Ablation

The grouped HGB ablation used the full 40-feature representation in the 2025 datasets. Removing soil-state history caused the largest degradation. For Station 34, full-model MAE increased from 10.8264, 16.1598, 17.3162, and 18.4084 to 21.5161, 22.0322, 22.0486, and 24.5644 after removing soil-state history at 1, 6, 12, and 24 h. For Station 35, the corresponding full-model values were 12.5140, 13.8953, 13.8011, and 13.7981, versus 27.0817, 26.4589, 26.5071, and 25.7135 without soil-state history.

6.4 Training-Window Sensitivity

Ridge 24-h forecasting was evaluated with initial training fractions of 40%, 50%, and 60% in the 2025 datasets. Skill versus persistence remained negative for all six station-window combinations.

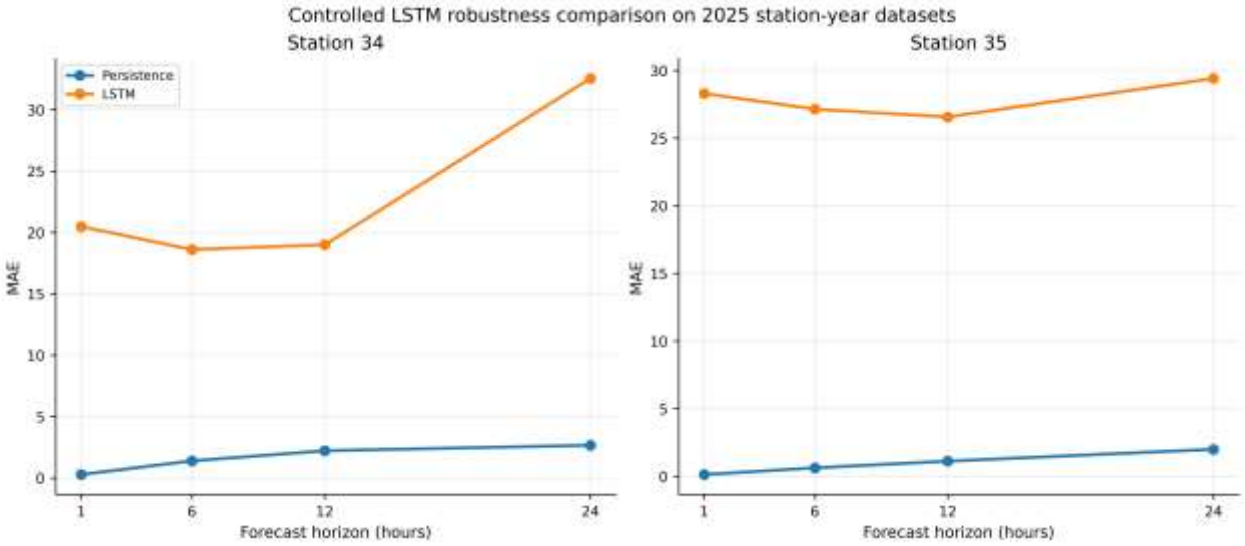


Fig. 13: Controlled LSTM robustness results for the 2025 station-year experiments.

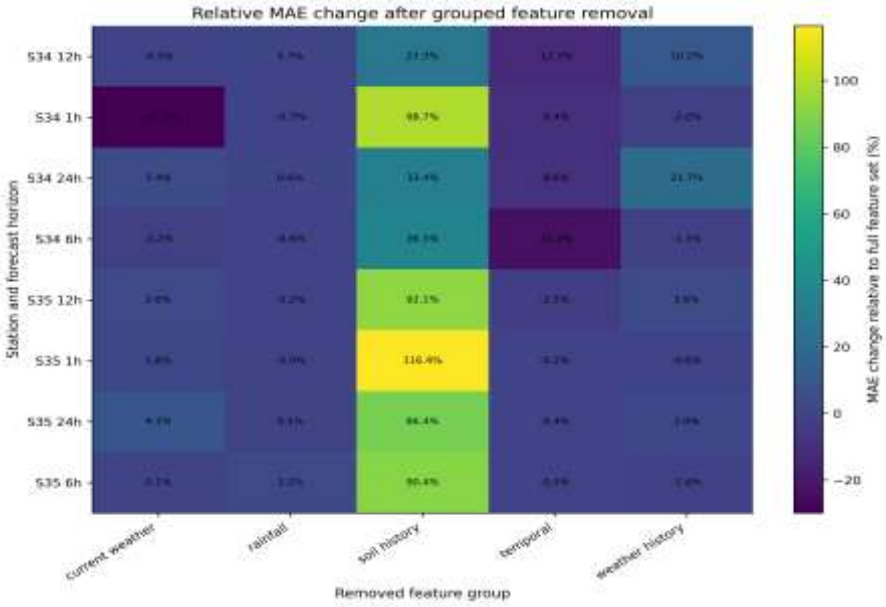


Fig. 14: Grouped feature-ablation MAE heatmap. Soil-state history is the most consequential feature group in the evaluated HGB robustness experiments, although no configuration surpasses persistence.

Natural Resources for Human Health
eISSN: 2583-1194

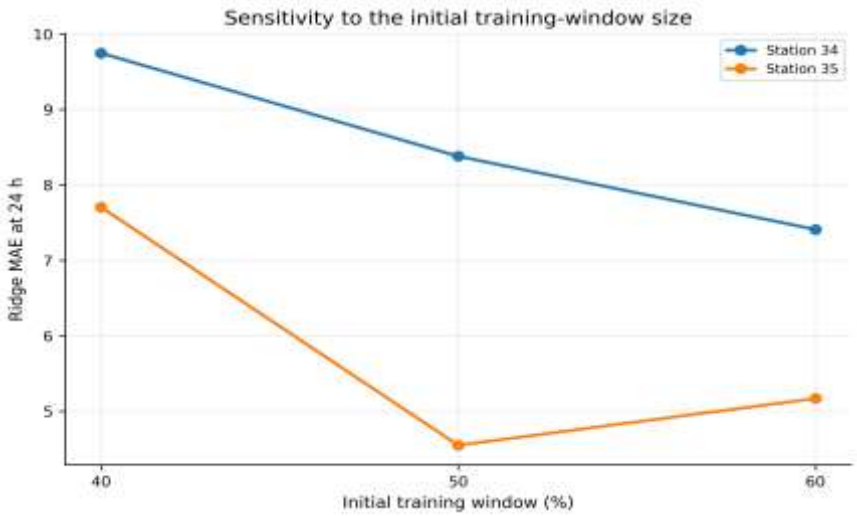


Fig. 15: Training-window sensitivity for 24-h Ridge forecasting. Negative skill remains across all tested initial training fractions.

6.5 Bootstrap Block-Length Sensitivity

For Ridge, all confidence intervals remained above zero across 12-, 24-, 48-, and 168-h blocks. At Station 34 and 24 h, the intervals were [5.1524, 6.6657], [4.8996, 6.9163], [4.6334, 7.2418], and [3.8940, 8.0262]. At Station 35, they were [1.9180, 3.2321], [1.7125, 3.4663], [1.4144, 3.8097], and [0.9013, 4.4571].

6.6 Complete Primary MAE Heatmap

6.7 Stationarity Diagnostics

Stationarity evidence was mixed at the sequence level. Station 34–2025 had 26.67% ADFstationary sequences, 20.00% KPSS-stationary sequences, 6.67% satisfying both criteria, and 60.00% satisfying neither. Station 35–2025 had 25.00%, 25.00%, 25.00%, and 75.00%, respectively.

Station 34–2024 had 33.33%, 0%, 0%, and 66.67%, while Station 35–2024 had 25.00%, 25.00%, 0%, and 50.00%. The results do not justify labeling an entire station-year as uniformly stationary or nonstationary.

7. INTEGRATED INTERPRETATION OF THE FINDINGS

7.1 Temporal Dependence and the Strength of Persistence

The consistently low persistence error and strong autocorrelation indicate that recent soilmoisture state contains substantial short-term predictive information. This provides a direct

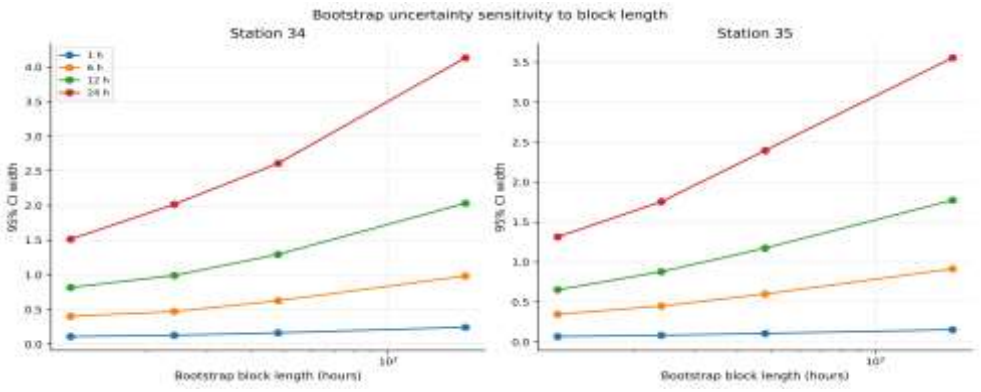


Fig. 16: Sensitivity of Ridge-minus-persistence confidence intervals to temporal bootstrap block length.

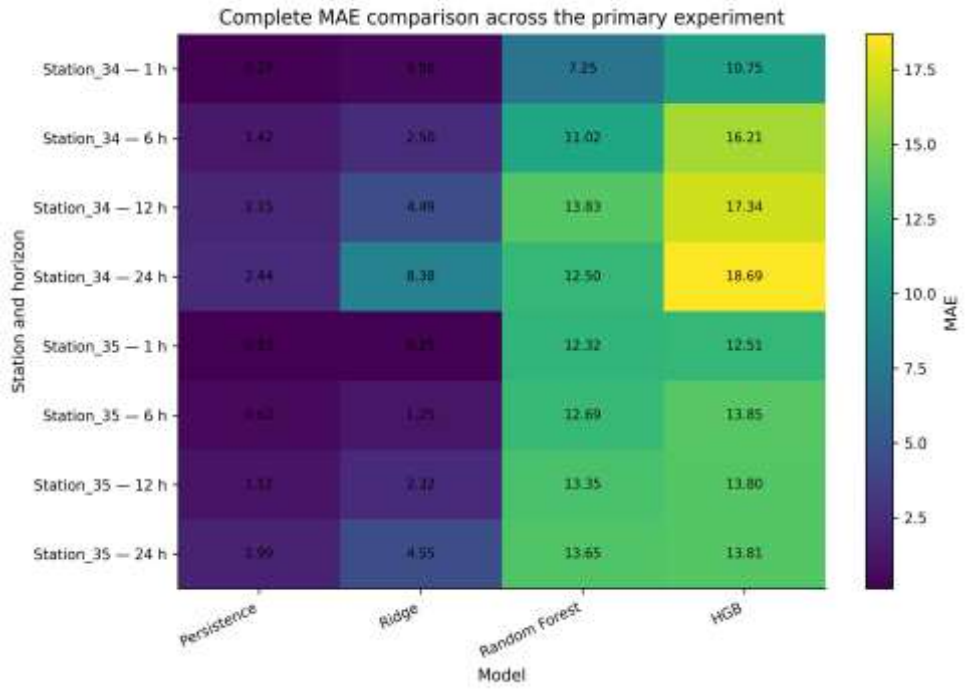


Fig. 17: Complete primary MAE heatmap across both stations, four horizons, and all evaluated models.

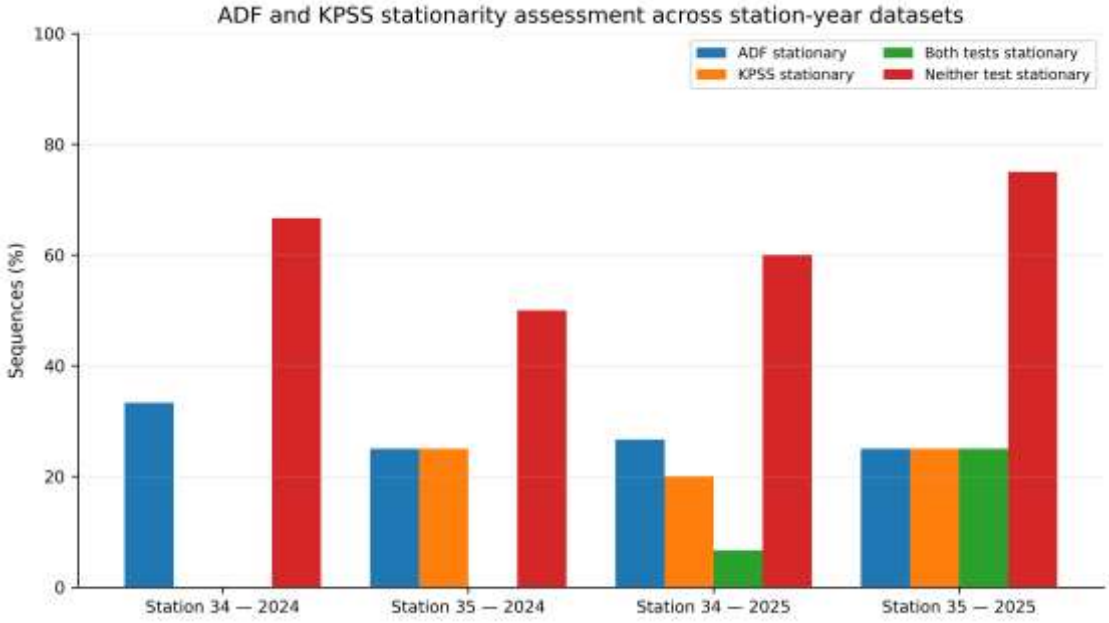


Fig. 18: Sequence-level ADF/KPSS stationarity summary for all four station-year datasets.

explanation for why high R^2 values from ML models do not automatically imply incremental value.

7.2 Forecast-Horizon Effects

The disadvantage of the best ML model increased with horizon. For Station 34, Ridge's relative MAE increase was 82.00%, 76.44%, 108.24%, and 242.95% at 1, 6, 12, and 24 h. For Station 35, it was 96.50%, 100.20%, 106.99%, and 128.08%. This pattern indicates that the relationship between model complexity and forecasting value is horizon-dependent.

7.3 Model Complexity

Random Forest and HGB performed substantially worse than Ridge and persistence. The result shows that greater nonlinear complexity did not automatically translate into improved chronological forecasting under this feature representation and dataset.

7.4 Feature Information

The feature-ablation results show that soil-state history is important to ML prediction. However, the full and ablated ML configurations remained below persistence. Thus, the result is not explained simply by the absence of useful lag information; the tested ML models had access to extensive historical and environmental features but still did not produce incremental accuracy.

7.5 Robustness of the Negative Result

The conclusion remained unchanged across alternative persistence baselines, the controlled LSTM, feature ablation, training-window choices, and bootstrap block lengths. This strengthens the interpretation that the primary result is not a single-model artifact.

Taken together, these checks address different potential sources of sensitivity: baseline choice, sequence-model capacity, dependence on feature groups, training-window size, bootstrap block length, and temporal stationarity. Their agreement supports the stability of the observed result within the tested design, while the geographic, feature, and model-family limitations described below remain applicable.

7.6 Implications for Smart Irrigation

The study evaluates soil-moisture forecasting rather than direct irrigation-event or irrigationvolume prediction because the public dataset does not contain a reliable explicit irrigation target. Soil-moisture forecasts can nevertheless form one component of an irrigation decision-support system. If a persistence forecast is more accurate than a tested ML model, adding ML complexity should be justified by demonstrable forecasting or downstream decision benefit in the target deployment environment.

7.7 Negative Results as a Benchmarking Contribution

The principal contribution is methodological as well as empirical. The study shows that a model can achieve high R^2 while still failing to improve upon a simple persistence forecast. Reporting strong baselines, chronological validation, and uncertainty intervals therefore provides a more informative assessment of whether ML adds value than reporting ML accuracy alone.

8. LIMITATIONS

1. Geographic and environmental scope: two monitoring stations from one public project do not represent the full diversity of agricultural soils, climates, crops, and sensor deployments.
2. No explicit irrigation target: the dataset does not provide a reliable irrigation-event or irrigation-volume target; therefore the study evaluates soil-moisture forecasting rather than direct irrigation prediction.
3. Model scope: the primary benchmark contains three classical ML models. The controlled LSTM is a robustness baseline, not an exhaustive deep-learning benchmark.
4. Feature scope: the feature set is restricted to variables available in the dataset. Soil properties, vegetation indices, remote sensing, irrigation records, and external numerical weather forecasts may provide information unavailable here.
5. Forecast horizon scope: only 1-, 6-, 12-, and 24-h horizons were evaluated; longer horizons and probabilistic forecasting may produce different conclusions.
6. Stationarity: ADF/KPSS evidence was mixed at sequence level, so no single stationary/nonstationary characterization is assigned to an entire station-year.
7. Representation and hyperparameters: the 40-feature representation and tested hyperparameters are one controlled design and do not exhaust all possible model configurations.

8. Generalization: conclusions are specific to the dataset, quality-control procedure, feature configuration, model families, hyperparameters, horizons, and chronological validation protocol.

9. CONCLUSION

This study investigated whether machine learning provides incremental predictive value beyond persistence for short-term soil-moisture forecasting across four station-year datasets. After quality control, 22,050 observations from two monitoring stations were retained. Ridge regression, Random Forest, and HGB were evaluated at 1-, 6-, 12-, and 24-h horizons using chronological expanding-window validation.

Persistence achieved the lowest MAE in all eight station-horizon combinations. Ridge was the strongest evaluated ML model but remained worse than persistence at every horizon. The mean relative MAE increase of the best ML model across the eight comparisons was 117.68%, with a median increase of 103.59%. All primary Ridge-minus-persistence bootstrap 95% confidence intervals remained above zero.

The robustness analyses were consistent with this result. Seasonal persistence was weaker than ordinary persistence, the controlled LSTM did not outperform persistence, soil-history ablation degraded ML but did not create an ML advantage, negative Ridge skill remained across training-window choices, and the Ridge disadvantage remained across bootstrap block lengths. Sequence-level stationarity tests showed mixed evidence rather than a uniform regime.

The principal conclusion is specific and empirical: under the evaluated dataset, feature configuration, model families, forecasting horizons, and chronological validation protocol, the tested ML approaches did not demonstrate incremental predictive value beyond persistence. This does not establish that ML is unsuitable for soil-moisture forecasting generally. Instead, it demonstrates the importance of strong persistence baselines, horizon-aware evaluation, chronological validation, and dependence-aware uncertainty analysis before interpreting model complexity as predictive improvement.

Future work should evaluate richer spatial and soil-property information, direct irrigation records, remote-sensing variables, external weather forecasts, probabilistic prediction, longer horizons, and more extensive sequence-based deep-learning architectures in larger multi-location datasets.

Data Availability

The environmental dataset analyzed in this study is publicly available through Zenodo under the title "AgriDataValue – Environmental Data Useful for Smart Irrigation" and DOI 10.5281/zenodo.18959383 [9]. No unverified code-repository URL is claimed.

Ethics Statement

The study uses environmental sensor observations and does not involve human participants, human-subject data, or animal experimentation. No human-subject ethical approval was therefore required for the analyses reported here, subject to the terms of the public dataset archive and license.

Code Availability

The original analysis code, generated numerical results, documentation, and citation metadata are publicly available in the reproducibility repository at https://github.com/pavanhari6999/soil-m oisture-forecasting. Raw third-party dataset files and publisher-copyrighted articles, figures, and tables are not redistributed.

CRediT Author Statement

Daggumati Pavan Hari Krishna: Conceptualization, Data curation, Methodology, Software, Formal analysis, Investigation, Visualization, Writing—original draft, Writing—review and editing.

Declaration of Competing Interest

The author declares that there are no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

The author acknowledges the AgriDataValue project and the dataset authors for providing the publicly accessible environmental observations used in this research.

Appendix A Complete Primary and Robustness Results

This appendix preserves the complete numerical outputs used in the manuscript. These tables are intentionally kept in the source so that no experimental result is silently omitted from the archival manuscript package. A target journal may move these tables to supplementary material if required by its page limit.

Table 8. Complete primary experimental results: MAE, RMSE, R2, ΔMAE, and bootstrap confidence intervals.

Station	h	Model	*P*MAE	MAE	RMSE	*R*2	ΔMAE	CI_L	CI_U	Winner
Station	h	Model	*P*MAE	MAE	RMSE	*R*2	ΔMAE	CI_L	CI_U	Winner
34	1	HGB	0.2737	10.7513	17.3184	0.4680	10.4776	8.2862	12.5172	P
34	1	RF	0.2737	7.2491	11.9199	0.7095	6.9754	5.4903	8.4462	P
34	1	Ridge	0.2737	0.4982	1.1749	0.9976	0.2245	0.1704	0.2775	P
34	6	HGB	1.4159	16.2134	22.3329	0.1036	14.7975	11.7724	17.5284	P
34	6	RF	1.4159	11.0179	16.5622	0.4013	9.6020	7.6739	11.6412	P
34	6	Ridge	1.4159	2.4982	4.4665	0.9646	1.0823	0.8938	1.2543	P
34	12	HGB	2.1543	17.3368	24.4632	-0.1680	15.1825	11.7404	18.4142	P
34	12	RF	2.1543	13.8255	19.9051	0.1567	11.6712	9.1518	14.1231	P
34	12	Ridge	2.1543	4.4861	7.0038	0.9075	2.3319	1.9480	2.6724	P
34	24	HGB	2.4435	18.6861	25.3449	-0.3761	16.2426	12.5307	19.9932	P
34	24	RF	2.4435	12.4970	17.9906	0.1729	10.0535	7.7842	12.3349	P
34	24	Ridge	2.4435	8.3801	10.4621	0.7499	5.9366	5.1566	6.6952	P
35	1	HGB	0.1293	12.5087	20.2181	0.5146	12.3794	8.7054	15.9192	P
35	1	RF	0.1293	12.3218	20.0930	0.5177	12.1925	8.5323	15.7335	P
35	1	Ridge	0.1293	0.2541	0.9917	0.9983	0.1248	0.0963	0.1544	P
35	6	HGB	0.6231	13.8525	21.0966	0.4878	13.2294	9.4828	16.7791	P
35	6	RF	0.6231	12.6895	20.4551	0.5039	12.0664	8.2975	15.5644	P
35	6	Ridge	0.6231	1.2475	3.3533	0.9825	0.6244	0.4481	0.8097	P
35	12	HGB	1.1197	13.7994	21.7084	0.4638	12.6797	8.9104	16.6525	P
35	12	RF	1.1197	13.3439	21.2403	0.4783	12.2241	8.5328	16.0778	P
35	12	Ridge	1.1197	2.3177	5.2331	0.9589	1.1984	0.8657	1.5636	P

35	24	HGB	1.9945	13.8036	21.6176	0.4547	11.8091	8.1043	15.6146	P
35	24	RF	1.9945	13.6477	21.5039	0.4632	11.6532	7.9305	15.4285	P
35	24	Ridge	1.9945	4.5491	8.7321	0.8871	2.5546	1.9148	3.2903	P

Appendix B Complete Fold-Level Primary Results

Table 9. Chronological fold-level primary results.

Station	Horizon/Fold	Model	MAE	RMSE	*R2*	P-MAE	P-RMSE	P-R^2
34	1/F1	Ridge	0.3293	0.4336	0.9994	0.0993	0.2198	0.9998
34	1/F1	RF	4.7432	10.9030	0.6101	0.0993	0.2198	0.9998
34	1/F1	HGB	5.4906	12.2414	0.5085	0.0993	0.2198	0.9998
34	1/F2	Ridge	0.6671	1.9163	0.9958	0.4482	1.7838	0.9964
34	1/F2	RF	9.7551	12.9368	0.8089	0.4482	1.7838	0.9964
34	1/F2	HGB	16.0120	22.3954	0.4274	0.4482	1.7838	0.9964
34	6/F1	Ridge	1.8026	2.2718	0.9812	0.4621	0.8512	0.9974
34	6/F1	RF	6.9289	14.9005	0.1913	0.4621	0.8512	0.9974
34	6/F1	HGB	6.2730	13.1801	0.3673	0.4621	0.8512	0.9974
34	6/F2	Ridge	3.1938	6.6613	0.9481	2.3697	6.6432	0.9484
34	6/F2	RF	15.1068	18.2239	0.6114	2.3697	6.6432	0.9484
34	6/F2	HGB	26.1539	31.4857	-0.1601	2.3697	6.6432	0.9484
34	12/F1	Ridge	3.3324	3.9970	0.9343	0.6785	1.2313	0.9938
34	12/F1	RF	6.9850	15.4314	0.0211	0.6785	1.2313	0.9938
34	12/F1	HGB	7.5455	15.8579	-0.0337	0.6785	1.2313	0.9938
34	12/F2	Ridge	5.6398	10.0107	0.8807	3.6300	9.6146	0.8899
34	12/F2	RF	20.6660	24.3788	0.2922	3.6300	9.6146	0.8899
34	12/F2	HGB	27.1281	33.0685	-0.3023	3.6300	9.6146	0.8899
34	24/F1	Ridge	5.9395	6.7094	0.7481	0.8144	1.4052	0.9889
34	24/F1	RF	5.8082	13.6360	-0.0407	0.8144	1.4052	0.9889
34	24/F1	HGB	6.8246	14.0189	-0.1000	0.8144	1.4052	0.9889
34	24/F2	Ridge	10.8207	14.2148	0.7517	4.0726	11.9371	0.8249
34	24/F2	RF	19.1858	22.3452	0.3865	4.0726	11.9371	0.8249

34	24/F2	HGB	30.5476	36.6708	-0.6522	4.0726	11.9371	0.8249
35	1/F1	Ridge	0.2506	0.9762	0.9971	0.1709	0.8914	0.9975
35	1/F1	RF	9.0215	14.4501	0.3548	0.1709	0.8914	0.9975
35	1/F1	HGB	9.0482	14.4498	0.3548	0.1709	0.8914	0.9975
35	1/F2	Ridge	0.2576	1.0072	0.9995	0.0877	0.9807	0.9995
35	1/F2	RF	15.6221	25.7360	0.6807	0.0877	0.9807	0.9995
35	1/F2	HGB	15.9692	25.9864	0.6744	0.0877	0.9807	0.9995
35	6/F1	Ridge	1.0151	3.0288	0.9716	0.8014	2.8176	0.9754
35	6/F1	RF	9.3302	14.5785	0.3424	0.8014	2.8176	0.9754
35	6/F1	HGB	9.3661	14.5586	0.3441	0.8014	2.8176	0.9754
35	6/F2	Ridge	1.4799	3.6779	0.9935	0.4449	3.1111	0.9953
35	6/F2	RF	16.0488	26.3317	0.6654	0.4449	3.1111	0.9953
35	6/F2	HGB	18.3389	27.6345	0.6315	0.4449	3.1111	0.9953
35	12/F1	Ridge	1.8274	4.5943	0.9346	1.3962	4.1602	0.9463
35	12/F1	RF	9.7514	14.7067	0.3295	1.3962	4.1602	0.9463

Continued on next page

Table 9. Chronological fold-level primary results (continued).

Station	Horizon/Fold	Model	MAE	RMSE	*R2*	P-MAE	P-RMSE	P-R^2
35	12/F1	HGB	9.7834	14.7611	0.3245	1.3962	4.1602	0.9463
35	12/F2	Ridge	2.8080	5.8719	0.9833	0.8432	4.5601	0.9899
35	12/F2	RF	16.9363	27.7739	0.6271	0.8432	4.5601	0.9899
35	12/F2	HGB	17.8155	28.6557	0.6030	0.8432	4.5601	0.9899
35	24/F1	Ridge	3.6811	7.5926	0.8216	2.5282	6.3624	0.8747
35	24/F1	RF	10.1451	14.9299	0.3101	2.5282	6.3624	0.8747
35	24/F1	HGB	10.4661	15.0956	0.2947	2.5282	6.3624	0.8747
35	24/F2	Ridge	5.4171	9.8716	0.9526	1.4608	6.6314	0.9786
35	24/F2	RF	17.1502	28.0778	0.6164	1.4608	6.6314	0.9786
35	24/F2	HGB	17.1411	28.1396	0.6147	1.4608	6.6314	0.9786

Natural Resources for Human Health
eISSN: 2583-1194

Appendix C Complete Robustness Tables

C-A. Seasonal Persistence

Table 10. Seasonal-persistence MAE in the 2025 robustness experiments.

Station	h	Ordinary	24-h seasonal	48-h seasonal	168-h seasonal
34	1	0.2762	3.7476	7.3633	20.2255
34	6	1.6063	4.6905	8.1658	20.4651
34	12	2.6095	5.7314	9.1448	20.7483
34	24	3.6720	7.1660	10.7151	21.3185
35	1	0.1282	2.0130	4.2026	12.1677
35	6	0.6157	2.4406	4.6336	12.3077
35	12	1.0972	2.9405	5.1084	12.3990
35	24	1.9191	3.8600	5.7509	12.2985

C-B. Feature Ablation

TABLE XI: Grouped HGB feature-ablation MAE for 2025 station-year datasets.

St.	h	Full	No rain	No soil hist.	No temporal	No current weather	No weather hist.
34	1	10.8264	10.7482	21.5161	9.8057	7.5920	10.6142
34	6	16.1598	16.0664	22.0322	12.4122	15.8056	15.9544
34	12	17.3162	17.4338	22.0486	15.1180	17.2337	19.0754
34	24	18.4084	18.5176	24.5644	16.8189	19.0307	22.4107
35	1	12.5140	12.5118	27.0817	12.4833	12.7393	12.4333
35	6	13.8953	14.2021	26.4589	13.8222	13.9099	13.6793
35	12	13.8011	13.7744	26.5071	13.4596	14.1655	14.2924
35	24	13.7981	13.8092	25.7135	13.7384	15.0814	14.0711

C-C. Training-Window Sensitivity

TABLE XII: Training-window sensitivity for 24-h Ridge forecasting in 2025.

Station	Train	MAE	RMSE	*R2*	P-MAE	Skill (%)
34	40%	9.7462	13.5399	0.7084	2.8578	-2.0781
34	50%	8.3801	10.4621	0.7499	2.4435	-3.9750

34	60%	7.4092	11.6993	0.8286	2.9674	-2.5480
35	40%	7.7053	11.5727	0.6439	1.8521	-3.0671
35	50%	4.5491	8.7321	0.8871	1.9945	-1.5822
35	60%	5.1705	9.4248	0.8909	2.0413	-1.8234

C-D. Bootstrap Block-Length Sensitivity

TABLE XIII: Ridge-minus-persistence bootstrap confidence intervals under alternative circular block lengths.

Station	h	12 h block	24 h block	48 h block	168 h block
34	1	[0.1725,0.2814]	[0.1631,0.2899]	[0.1454,0.3067]	[0.1071,0.3493]
34	6	[0.8803,1.2807]	[0.8492,1.3193]	[0.7722,1.3969]	[0.5986,1.5810]
34	12	[1.9217,2.7385]	[1.8367,2.8264]	[1.6841,2.9779]	[1.3513,3.3830]
34	24	[5.1524,6.6657]	[4.8996,6.9163]	[4.6334,7.2418]	[3.8940,8.0262]
35	1	[0.0937,0.1602]	[0.0867,0.1684]	[0.0757,0.1801]	[0.0543,0.2061]
35	6	[0.4547,0.7996]	[0.4085,0.8561]	[0.3343,0.9330]	[0.1924,1.1057]
35	12	[0.8771,1.5283]	[0.7736,1.6514]	[0.6319,1.8051]	[0.3620,2.1344]
35	24	[1.9180,3.2321]	[1.7125,3.4663]	[1.4144,3.8097]	[0.9013,4.4571]

C-E. ADF/KPSS Stationarity

TABLE XIV: Sequence-level ADF/KPSS stationarity summary.

Dataset	ADF stationary	KPSS stationary	Both	Neither
Station 34–2024	33.33%	0.00%	0.00%	66.67%
Station 35–2024	25.00%	25.00%	0.00%	50.00%
Station 34–2025	26.67%	20.00%	6.67%	60.00%
Station 35–2025	25.00%	25.00%	25.00%	75.00%

C-F. MAE Bootstrap Confidence Intervals

TABLE XV: Fold-level MAE bootstrap confidence intervals for persistence, Ridge, and their difference.

St.	h	Fold	P MAE CI	Ridge MAE CI	ΔMAE CI
34	1	1	[0.0726,0.1287]	[0.2971,0.3663]	[0.2123,0.2482]

34	1	2	[0.3088,0.6133]	[0.5425,0.8224]	[0.1215,0.3238]
34	6	1	[0.3418,0.6023]	[1.6169,1.9860]	[1.1697,1.5164]
34	6	2	[1.6554,3.2341]	[2.6149,3.9549]	[0.5440,1.1437]
34	12	1	[0.4753,0.8851]	[2.9917,3.6482]	[2.3118,3.0344]
34	12	2	[2.3001,5.2859]	[4.5457,7.1808]	[1.3818,2.6411]
34	24	1	[0.5545,1.0964]	[5.3601,6.5269]	[4.4241,5.8325]
34	24	2	[2.3113,6.2341]	[9.2481,12.6659]	[5.4706,7.9740]
35	1	1	[0.1041,0.2656]	[0.1657,0.3681]	[0.0447,0.1240]
35	1	2	[0.0365,0.1811]	[0.1829,0.3719]	[0.1262,0.2119]
35	6	1	[0.4973,1.2054]	[0.6180,1.5532]	[0.0500,0.3857]
35	6	2	[0.1535,0.9885]	[1.0022,2.2099]	[0.7547,1.3353]
35	12	1	[0.8026,2.3702]	[1.0729,2.8965]	[0.1158,0.7521]
35	12	2	[0.2443,1.9363]	[1.8668,4.1846]	[1.4564,2.5118]
35	24	1	[1.2350,4.5699]	[2.0891,5.9090]	[0.4320,1.8755]
35	24	2	[0.3227,3.5095]	[3.5317,8.0487]	[2.9302,5.0456]

Appendix D Figure Inventory

The manuscript contains 18 figure elements. The 17 empirical figures correspond to the verified final figure package: Station 34 MAE; Station 35 MAE; RMSE; R^2; relative MAE increase; bootstrap ΔMAE; QC observation counts; soil-moisture distributions; temporal coverage; autocorrelation; seasonal persistence; feature-ablation heatmap; training-window sensitivity; bootstrap block-length sensitivity; complete MAE heatmap; LSTM robustness; and ADF/KPSS stationarity. Figure 18 is the QC retention pie chart generated directly from the verified totals.

REFERENCES

[1] I. Ivanova, "Soil moisture forecasting from sensors-based soil moisture, weather and irrigation observations: A systematic review," *Smart Agricultural Technology*, vol. 10, Art. no. 100692, 2025, doi: 10.1016/j.atech.2024.100692.

[2] M. B. Islam, A. Guerrieri, R. Gravina, D. T. Delaney, and G. Fortino, "From traditional machine learning to fine-tuning large language models: A review for sensors-based soil moisture forecasting," *Sensors*, vol. 25, no. 22, Art. no. 6903, 2025, doi: 10.3390/s25226903.

[3] R. Togneri, D. F. dos Santos, G. Camponogara, H. Nagano, G. Custodio, R. C. Prati, S. Fernandes, and C. Kamienski, "Soil moisture forecast for smart irrigation: The primetime for machine learning," *Expert Systems with Applications*, vol. 207, Art. no. 117653, 2022, doi: 10.1016/j.eswa.2022.117653.

[4] G. Custodio and R. C. Prati, "Comparing modern and traditional modeling methods for predicting soil moisture in IoT-based´ irrigation systems," *Smart Agricultural Technology*, vol. 7, Art. no. 100397, 2024, doi: 10.1016/j.atech.2024.100397.

[5] W. Zheng, K. Zheng, L. Gao, L. Zhangzhong, R. Lan, L. Xu, and J. Yu, "GRU–Transformer: A novel hybrid model for predicting soil moisture content in root zones," *Agronomy*, vol. 14, no. 3, Art. no. 432, 2024, doi: 10.3390/agronomy14030432.

[6] G. W. Tefera, R. L. Ray, R. Jackson, B. Deegala, and O. Akenzua, "Integrating machine learning models with ground sensors to enhance soil moisture prediction in agroecosystems of Texas," *Computers and Electronics in Agriculture*, vol. 242, Art. no. 111358, 2026, doi: 10.1016/j.compag.2025.111358.

[7] Q. Su, Y. Lin, G. C. Rains, and J. Mohammadpour Velni, "Transferable field-scale daily soil moisture forecasting in cotton with a minimal-input LSTM for irrigation scheduling," *Smart Agricultural Technology*, vol. 14, Art. no. 102178, 2026, doi: 10.1016/j.atech.2026.102178.

[8] M. Lamichhane, S. Mehan, and K. R. Mankin, "Soil moisture prediction using remote sensing and machine learning algorithms: A review on progress, challenges, and opportunities," *Remote Sensing*, vol. 17, no. 14, Art. no. 2397, 2025, doi: 10.3390/rs17142397.

[9] E. Ilie and N. Arvanitis, "AgriDataValue – Environmental Data Useful for Smart Irrigation," Zenodo, 2026, doi: 10.5281/zenodo.18959383.

[10] D. N. Politis and H. White, "Automatic block-length selection for the dependent bootstrap," *Econometric Reviews*, vol. 23, no. 1, pp. 53–70, 2004, doi: 10.1081/ETC-120028836.

[11] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American Statistical Association*, vol. 74, no. 366, pp. 427–431, 1979.

[12] D. Kwiatkowski, P. C. B. Phillips, P. Schmidt, and Y. Shin, "Testing the null hypothesis of stationarity against the alternative of a unit root," *Journal of Econometrics*, vol. 54, nos. 1–3, pp. 159–178, 1992, doi: 10.1016/0304-4076(92)90104-Y.

[13] X. Wang, G. Corzo, H. Lu, S. Zhou, K. Mao, Y. Zhu, S. Duarte, M. Liu, and J. Su, "Sub-seasonal soil moisture anomaly¨ forecasting using combinations of deep learning, based on the reanalysis soil moisture records," *Agricultural Water Management*, vol. 295, Art. no. 108772, 2024, doi: 10.1016/j.agwat.2024.108772.

[14] F. T. Teshome, H. K. Bayabil, B. Schaffer, Y. Ampatzidis, and G. Hoogenboom, "Improving soil moisture prediction with deep learning and machine learning models," *Computers and Electronics in Agriculture*, vol. 226, Art. no. 109414, 2024, doi: 10.1016/j.compag.2024.109414.

[15] M. Blahov´ a, M. Fischer, M. Pod´ ebradsk˘ a, P.´ St˘ ep˘ anek, J. Balek, P. Zahradn´ ´ıcek, L. Kudl˘ a´ckov˘ a, Z.´ Zalud, and M. Trnka,˘ "Testing the reliability of soil moisture forecast for its use in agriculture," *Agricultural Water Management*, vol. 304, Art. no. 109073, 2024, doi: 10.1016/j.agwat.2024.109073.