Received: 3 October 2025 | Accepted: 13 January 2026

DOI: 10.1002/vzj2.70077

Vadose Zone Journal

ORIGINAL ARTICLE

Special Section: Advances in Monitoring Soil Water Content

Unraveling soil moisture dynamics with dual-scale interpretable machine learning: Cover cropping and irrigation insights in semi-arid agriculture

Huichao Yin | **Prakriti Bista** | **Rajan Ghimire** | **Hui Yang** | **Kenneth C. Carroll**

Plant & Environmental Sciences Department, New Mexico State University, Las Cruces, New Mexico, USA

Correspondence
Kenneth C. Carroll, Plant & Environmental Sciences Department, New Mexico State University, Las Cruces, NM 88003, USA.
Email: kccarr@nmsu.edu

Assigned to Associate Editor Todd Caldwell.

Funding information
Savannah River National Laboratory, Grant/Award Number: TOA 800002114; National Institute of Food and Agriculture, Grant/Award Numbers: 2022-67019-36106, 2024-68017-42789

Abstract

Soil moisture prediction remains a major challenge in arid and semi-arid agroecosystems, where highly variable rainfall, intermittent irrigation, and heterogeneous soil properties create strong nonlinearities in soil-water dynamics. Reliable forecasts and interpretable models are needed for optimizing irrigation scheduling and improving water-use efficiency under such conditions. The purpose was development of a novel, dual-resolution framework that integrates machine learning (ML) and deep learning (DL) approaches with explainable artificial intelligence to evaluate soil moisture dynamics under different cover crop treatments. A 2-year field experiment was conducted in Clovis, NM, where soil moisture, temperature, and meteorological variables were continuously monitored across five treatments: fallow, pea, oat, pea–oat mixture, and a six-species mixture. ML models (random forest, light gradient boosting machine [LGBM], and eXtreme gradient boosting) and DL models, such as long short-term memory and Transformer models, were trained on both daily and 5-min datasets. SHapley Additive exPlanations (SHAP) indicates that features from the three most recent days had the strongest influence on daily predictions, informing a 72-h lookback for short-term modeling. LGBM achieved the highest accuracy, while Transformer reduced temporal lag in high-frequency predictions. SHAP analysis further identified treatment-specific sensitivities and scale-dependent relationships between environmental variables and soil moisture content, with historical soil moisture and total water input dominating daily predictions, and solar radiation, temperature, and recent water events prevailing at finer scales. This frame-

Abbreviations: DL, deep learning; LGBM, light gradient boosting machine; LSTM, long short-term memory; ML, machine learning; MLP, multi-layer perceptron; MSE, mean squared error; RF, random forest; RMSE, root mean squared error; SHAP, SHapley Additive exPlanations; SM, soil moisture; SR, solar radiation; SSM, six-species mixture; ST, soil temperature; TCN, temporal convolutional network; XAI, explainable AI; XGBoost, eXtreme gradient boosting.

work improves both predictive accuracy and interpretability of soil water dynamics, supporting data-driven irrigation and crop management in arid and semi-arid regions.

Plain Language Summary

To support decision-making, we developed a data-driven approach to predict how soil moisture changes over time and identify the main influencing factors. We tested both machine learning and deep learning models on data collected at two time scales: daily averages and high frequency (every 5 min). We used an interpretation tool, SHapley Additive exPlanations, to understand which factors were most important for the predictions. Soil moisture, irrigation, and precipitation over the past 3 days were the strongest predictors of moisture content at the daily scale, while weather factors like solar radiation and temperature became more important at finer time scales. Within 72 h of water input, cover crops (especially oat and the six-species mixture) showed greater water infiltration and uptake than fallow plots. These insights can help farmers optimize irrigation scheduling, choose cover crops that conserve water, and improve drought resilience in arid and semi-arid regions.

1 | INTRODUCTION

Soil moisture (SM) is a critical variable governing water availability, crop productivity, and ecosystem resilience (Sharma et al., 2020) in arid and semi-arid agroecosystems (Ghimire et al., 2023). In the southern High Plains of the United States, including eastern New Mexico and west Texas, limited and highly variable precipitation (Adebayo et al., 2025) combined with declining groundwater resources in the Ogallala Aquifer make efficient soil water management an urgent challenge for sustaining agriculture (Nilahyane et al., 2023). Soils in these regions are typically low in organic matter, slightly alkaline, and limited in water-holding capacity, further amplifying the risk of crop failure under variable rainfall and irrigation regimes (Thapa et al., 2022). There is a need for development of interpretable and predictive tools that can anticipate SM fluctuations and quantitatively evaluate influencing factors, and these would become essential for optimizing irrigation scheduling and guiding adaptive management, especially for agriculture in semi-arid regions (Zhang et al., 2025).

Cover cropping is increasingly promoted as a strategy to enhance soil health (Thapa et al., 2022) and water use efficiency (Chen et al., 2024) by improving aggregation, increasing soil organic carbon (Cheng et al., 2024), and reducing evaporation (Chen et al., 2024). However, its effects on soil water dynamics are often complex, particularly in water-limited environments (Zhang et al., 2024). Certain cover crops may conserve soil water and increase infiltration, whereas others may deplete available water due to competition (Furtak & Wolinska, 2023; Peng et al., 2024). Compost and organic amendments can also alter soil water retention, yet their interactions with cover crops in semi-arid conditions remain insufficiently characterized (Adebayo et al., 2025). We still need to improve our understanding of how these management practices influence SM across timescales, which is thus vital for improving water resilience in dryland and irrigated systems (Qin et al., 2024; Qiu et al., 2024).

Advances in data science now enable new ways to model SM dynamics (Sharma et al., 2020) by capturing nonlinear relationships between environmental variables. Machine learning (ML) and deep learning (DL) models have demonstrated strong performance in hydrological (Dai et al., 2025; Zhan et al., 2024, 2025) and agricultural forecasting (Yang et al., 2020), often outperforming traditional regression approaches (Teshome et al., 2024). Ensemble boosting methods such as eXtreme gradient boosting (XGBoost) and light gradient boosting machine (LGBM) can efficiently capture nonlinear feature interactions (Teshome et al., 2024), while sequential DL models including artificial neural network (Kornelsen & Coulibaly, 2014), long short-term memory (LSTM) (Bakhshian et al., 2025; Ley et al., 2024), and Transformer (Wang, Shi, et al., 2025) can leverage high-frequency time series to reduce temporal lag and identify complex dependencies. Yet, applications of these models to field-scale, multi-treatment experiments with multiple temporal resolutions remain limited, especially in semi-arid cropping systems influenced by both irrigation and cover crop practices.

For such predictive approaches to inform management, interpretability is crucial. Black-box models may achieve high accuracy but provide limited insight into the processes they represent (Ekanayake et al., 2022). Explainable AI (XAI) tools (Ley et al., 2024), particularly SHapley Additive exPla-

nations (SHAP), offer a transparent framework to quantify feature contributions and assess how weather, soil, and management factors influence model predictions (Nikraftar et al., 2025). Rather than inferring direct causality, SHAP provides data-driven insight into model sensitivities that can complement agronomic knowledge and guide decision-making in irrigation and soil management (Ley et al., 2024). To date, there has been very little research that examines DL model performance and XAI prediction control evaluations for SM, especially considering complex agricultural cropping systems. Teshome et al. (2024) did provide the only directly comparable example of applying ML and DL models for SM prediction using field data; however, their analysis was limited to a sub-hourly scale and reported only static feature rankings.

To evaluate the potential of multiple time resolution approaches, the purpose of this investigation was to develop a dual-resolution, SHAP-integrated framework that jointly assesses predictive accuracy and feature influence across daily and sub-daily timescales under diverse cover crop treatments in a semi-arid agroecosystem. We developed a dual-resolution interpretable modeling framework that uniquely integrates ML and DL approaches with SHAP to (i) predict SM at both daily and sub-daily temporal resolutions, (ii) quantify treatment-specific feature influence across timescales, and (iii) link data-driven model sensitivities to process-level understanding of irrigation and cover crop impacts in semi-arid systems. Unlike previous studies that focus solely on prediction or interpretation at a single scale, part of the novelty of this study bridges temporal resolutions to enhance both the accuracy and interpretability of SM dynamics. The framework thereby advances XAI applications in soil–water management by explicitly linking feature-influence analysis to experimental field conditions.

Core Ideas

- A dual-resolution framework integrates daily and 5-min data to enhance soil moisture prediction.
- Scale-dependent drivers of soil moisture include historical patterns, water inputs, and weather variables.
- SHapley Additive exPlanations (SHAP) analysis reveals cover crop sensitivities and identifies critical thresholds for water input.

2 | STUDY SITE AND DATASET DESCRIPTION

The present study builds on a 2-year field experiment in which SM, temperature, and meteorological variables were continuously monitored across five treatments: fallow (i.e., no cover), pea, oat, pea–oat mixture, and a six-species mixture (SSM). The dataset analyzed in this study was obtained from a field experiment conducted at the New Mexico State University Agricultural Science Center near Clovis, NM (34°35′ N, 103°12′ W; elevation 1368 m), spanning the period from July 16, 2022, to June 23, 2024. The experimental site is underlain by Olton clay loam soils (Aridic Paleustolls), characterized by low organic matter, slightly alkaline pH, and limited natural fertility (Ghimire et al., 2019). The local climate is semi-arid, with a mean annual temperature of 14.1°C and an average annual precipitation of about 437 mm, roughly 70% of which falls during the May–September growing season. A randomized complete block design was implemented with five cover crop treatments: (i) Fallow (no cover crop, F), (ii) pea (*Pisum sativum* L.), (iii) oat (*Avena sativa* L.), (iv) a pea–oat mixture (PO), and (v) an SSM comprising pea, oat, canola (*Brassica napus* L.), barley (*Hordeum vulgare* L.), hairy vetch (*Vicia villosa* Roth), and forage radish (*Raphanus sativus* L.). Each plot measured 18 m × 12 m. The experiment was conducted across five management phases: (i) sorghum (July 16–November 3, 2022), (ii) first fallow (November 6, 2022–February 22, 2023), (iii) cover crop (March 10–May 22, 2023), (iv) second fallow (May 23–October 8, 2023), and (v) wheat (October 27, 2023–June 23, 2024).

Soil volumetric water content (SM) and soil temperature (ST) were continuously monitored at a 5-cm depth in each plot using wireless sensors (Meter ECH2O, EC-5, Onset Computer Corporation). The data were collected at 5-min intervals, transmitted to an RX3000 remote monitoring station (Onset Computer Corporation), and uploaded to the HOBOlink cloud platform. The PO cropping data were discarded due to a large number of missing values caused by sensor malfunction. Weather variables, including average precipitation, air temperature (Tmean), humidity (RHmean), solar radiation (SR), and wind speed (Wmean), were also recorded with the identical frequency on site. Sensors were temporarily removed during field operations, such as planting and harvest, resulting in short gaps in the records. These gaps were filled using linear interpolation to ensure continuity. Irrigation and field management activities were documented throughout the experiment. Total water input was calculated using the same 5-min interval as the sum of irrigation and rainfall, serving as a key predictor variable for modeling SM dynamics and quantifying the external hydrological inputs driving soil water balance in the experimental plots. For data preprocessing, SM values were constrained to the range 0.02–0.557 $m^3\ m^{-3}$. The upper limit reflects the sensor's measurement capacity specified by the manufacturer, while the lower limit was set to exclude nonphysical values outside agronomically realistic soil conditions.

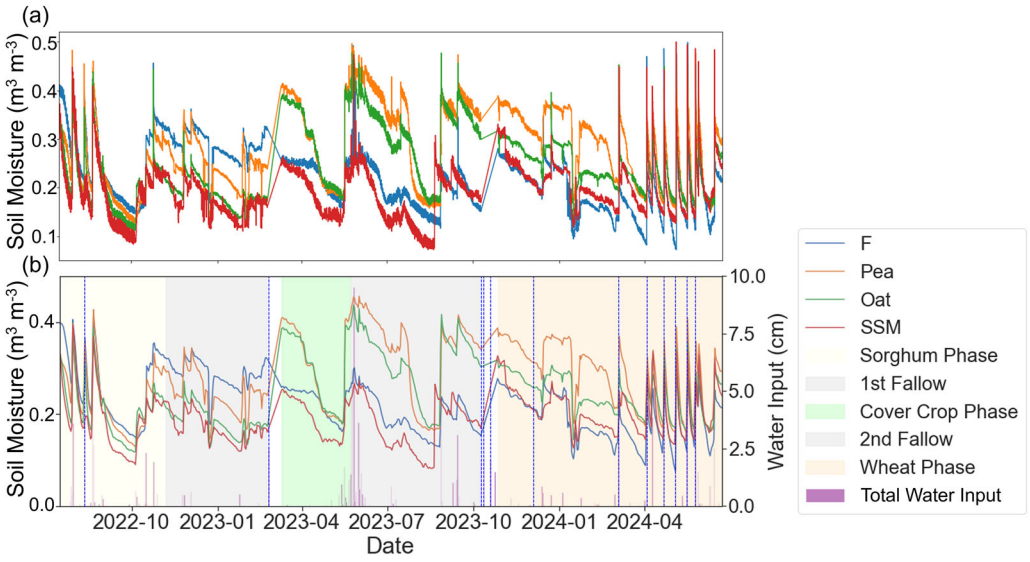

FIGURE 1 Temporal variation of the soil moisture in each cover crop plot in the field throughout the period of observation. The 5-min interval data (a) are resampled to daily mean values (b). Blue-dashed vertical lines indicate irrigation event dates (Table 1). Soil moisture data are unavailable during field operations and, therefore, linearly interpolated in unshaded periods (b).

TABLE 1 Irrigation and field operation conditions over time.

Date	Irrigation (cm)	Agriculture activity
August 11, 2022	4.572	Irrigation to sorghum
February 24, 2023	2.54	Cover crop planting + irrigation
October 9, 2023	2.032	Irrigation to winter wheat
October 12, 2023	2.032	Irrigation to winter wheat
October 19, 2023	2.032	Irrigation to winter wheat
March 4, 2024	3.048	Irrigation to winter wheat
April 3, 2024	2.54	Irrigation to winter wheat
April 21, 2024	3.81	Irrigation to winter wheat
May 4, 2024	4.572	Irrigation to winter wheat
May 16, 2024	4.572	Irrigation to winter wheat
May 25, 2024	4.572	Irrigation to winter wheat

Figure 1 and Table 1 illustrate the temporal dynamics of SM in each cover crop treatment across these phases. Panel (a) shows the original 5-min interval measurements, while panel (b) displays daily averages resampled from the high-frequency data. In total, 197,827 samples for the 5-min dataset and 709 for the daily dataset were obtained. The shaded backgrounds indicate crop phases, and unshaded periods correspond to sensor removal during field operations. Great fluctuations in the SM variations can be frequently seen upon major water inputs, including rainfall and irrigation, which is common for semi-arid regions. Curves of various cover crops demonstrate consistent variation trends despite value discrepancies. The resampled daily data were used to investigate mid- to long-term prediction of SM and feature importance, while the original high-frequency data were used for short-term prediction and fine-scale feature influence analysis. The dataset integrates high-frequency SM measurements with detailed records of weather conditions, irrigation, and management activities, providing a continuous, well-documented time series across multiple crop phases. This richness and resolution make it particularly well-suited for ML-based prediction and feature-influence analysis, enabling robust evaluation of both short-term and seasonal drivers of SM variability.

The statistical features of the daily dataset are presented in Figure 2. Figure 2a shows that median and quartile SM varied across treatments, with the Pea plot maintaining the highest values above 0.3 m^3/m^3 on average, the SSM plot the lowest below 0.2 m^3/m^3 on average, and the oat and fallow (F) plots falling in between. Despite low average values, SSM shows higher variability, with numerous outliers near 0.4, likely due to the mixed rooting patterns and different phe-

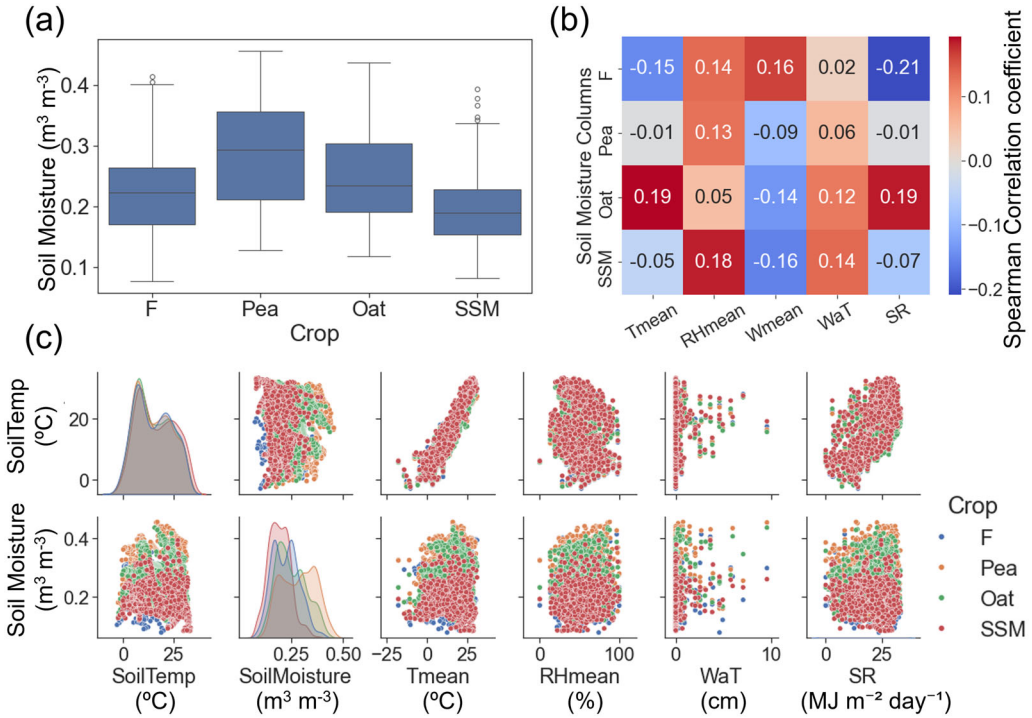

FIGURE 2 Soil moisture data statistical analyses including boxplot (a), Spearman correlation (b), heatmap with weather variables and pairplot (c), with soil temperature and weather variables. SSM, six-species mixture.

nologies of the species in the mixture (Cui et al., 2022; Luo et al., 2021). Such outliers of larger values are also presented in F and may correspond to episodic rainfall or irrigation events that produced temporary spikes in SM (Figure 1). Figure 2b demonstrates that the correlations between SM and other variables are generally weak, with most absolute correlation coefficients ˂0.2, whereas humidity and the water inputs are consistently and positively associated with SM. The strength of correlation varies by treatment: SSM plots show the strongest positive association with RHmean (0.18), and fallow plots display the most pronounced negative correlation with SR (−0.21), consistent with the lack of canopy shading. Figure 2c illustrates these interactions in greater detail. For ST, all treatments have generally overlapped data points, indicating similar correlations with other variables. For SM, despite weak correlations with other variables, treatment differences are apparent, demonstrated by diverse kernel density distributions, and noticeable clustering of vertical "bands" for correlations to each variable, which are pea, oat, SSM, and fallow, from top to bottom. This indicates that pea crop/soil systems can maintain higher soil water content than other treatments and fallow. Together, these statistical summaries demonstrate clear treatment effects and consistent meteorological controls, underscoring the dataset's richness and suitability for ML/DL-based prediction and feature influence analysis.

3 | METHODOLOGY

Building on a multivariate time-series dataset containing features of SM, ST, Tmean, RHmean, Wmean, and SR for cover crops (i.e., F, pea, oat, and SSM), a two-stage modeling framework (i.e., two temporal scales) for SM prediction and feature interpretation was developed. All water input to the soil (WaT; i.e., the combination of precipitation and irrigation) was included as a time-series dataset. SM prediction in this setting is inherently challenging because rainfall and irrigation events occur stochastically (Figure 1), driving abrupt fluctuations in soil water content. These short-lived, high-magnitude changes introduce strong nonlinearities and temporal irregularities, complicating model training and forecasting. To address this challenge, the workflow was designed to leverage both daily-aggregated data and high-frequency (5 min) measurements, applying a combination of ML and DL models, complemented by SHAP-based XAI analysis. The workflow proceeded in two phases (Figure 3). In the first phase, daily data were used to predict mid- to long-term SM. Each model was trained to predict SM for the following day using information from the past 14 days. Daily data were chosen for this stage because they capture coarser-scale water dynamics and yield fewer samples, both of which constrain training and emphasize model robustness across crop treatments. In the second phase, we extended the

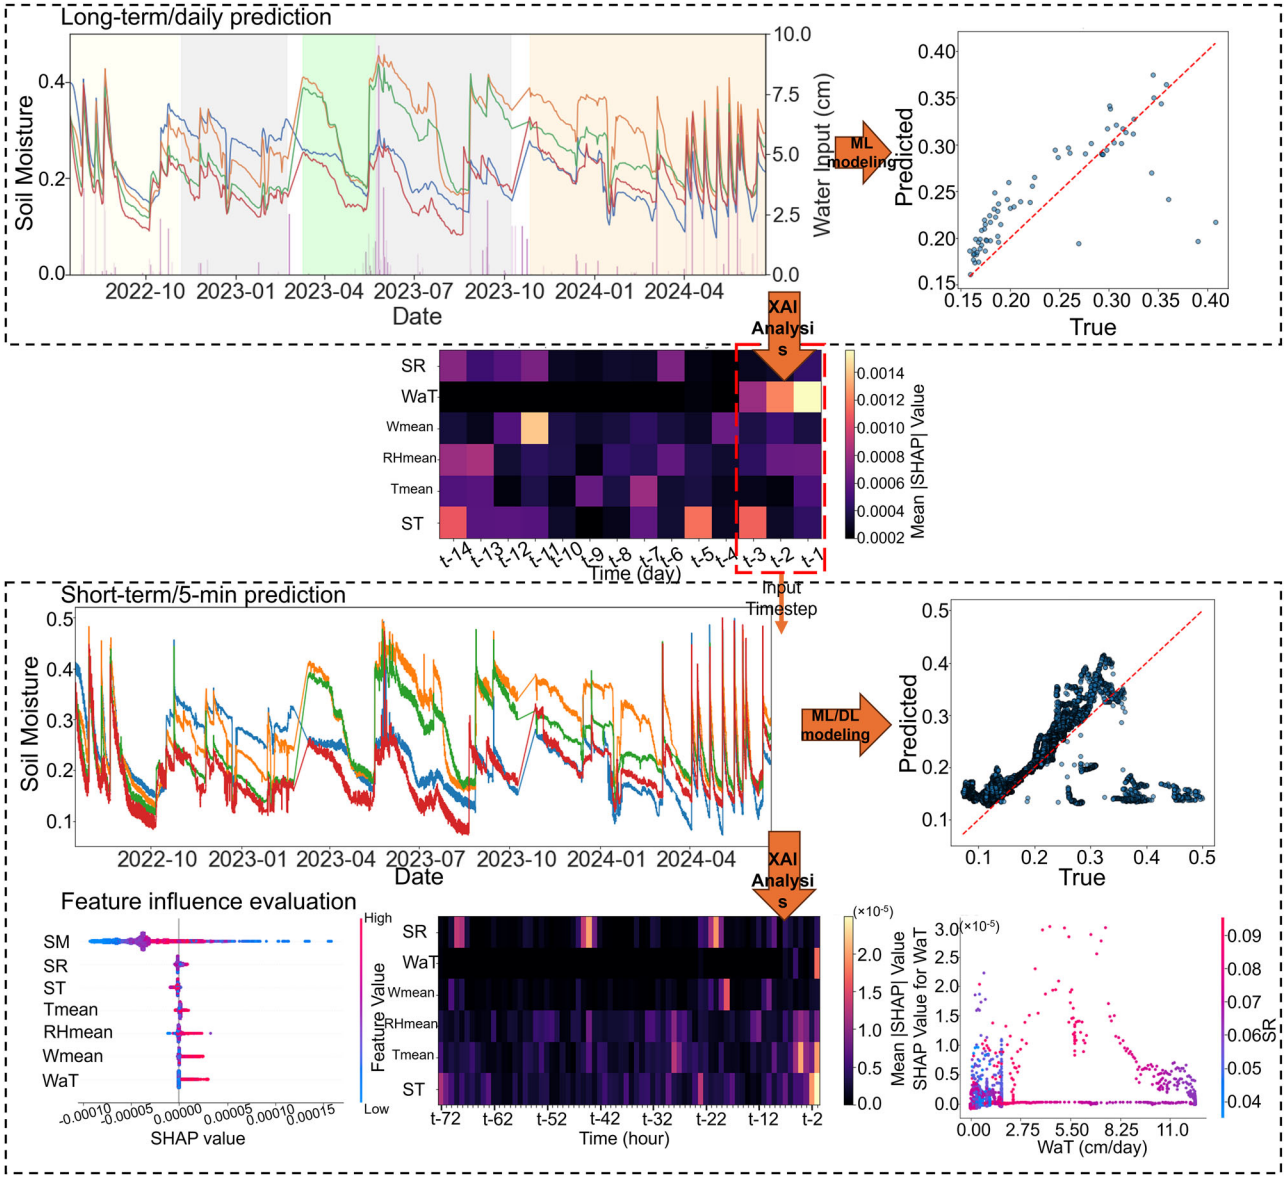

FIGURE 3 Schematic diagram of the dual-resolution soil moisture prediction and input feature evaluation framework. Phase 1 (top dashed box) uses daily data with a 14-day lookback to predict the following day and assess treatment-specific dynamics with the SHapley Additive exPlanations (SHAP) feature and timestep attribution. Phase 2 (top dashed box) applies a SHAP-informed 72-h lookback on 5-min data to predict the next 8 h, benchmarking tree-based models against deep learning (DL) architectures while evaluating accuracy, lag, and feature influences. ML, machine learning; SM, soil moisture; SR, solar radiation; ST, soil temperature; XAI, explainable AI.

analysis to high-resolution (5 min) data for short-term prediction tasks to enhance predictive accuracy, temporal alignment between predicted and observed values (lag analysis), and the robustness of predictions under sudden fluctuations in water input. The best-performing daily models (LGBM and XGBoost, consistent with Teshome et al., 2024) were selected as benchmarks and compared with advanced DL models. The lookback window for the short-term task was specified based on SHAP-derived timestep importance from the daily models, using the past 72 h to predict SM for the next 8 h.

To ensure robust comparisons and mitigate potential feature interdependence, SM for each cover crop treatment was modeled independently and comparatively. This design enables (i) assessment of model training stability and generalization across treatments, (ii) diagnosis of feature importance and model interpretability using SHAP analysis, and (iii) evaluation of cover crop influences on SM dynamics under shared meteorological conditions. Modeling each treatment separately also minimizes the effect of high correlations among SM variables across different cover crops, which could otherwise bias the learning process. In contrast, correlations between SM and weather-related predictors remain low (Figure 2), indicating that the predictive relationships are largely driven by distinct environmental inputs rather

than redundant temporal patterns. With SHAP, we applied global and local explanations using aggregated feature importance, timestep attribution, and visualization techniques such as beeswarm plots, dependency plots, and temporal heatmaps. These analyses provided insight into how meteorological and water input factors (i.e., precipitation, irrigation) interact to shape SM variability across timescales.

3.1 | ML models

Tree-based ensemble models were selected as the baseline for the SM prediction task because of their ability to capture nonlinear relationships, handle high-dimensional features, and remain robust to multicollinearity (Demir & Sahin, 2023; Gu et al., 2023; Wang, Gao et al., 2025). These models are also well-compatible with SHAP analysis because SHAP values can be computed exactly for tree ensembles using the TreeSHAP algorithm, which leverages the additive tree structure to provide efficient, consistent, and locally accurate feature attributions, thereby facilitating interpretability and input feature evaluation (Li et al., 2024).

1. XGBoost: XGBoost is a gradient boosting algorithm designed for scalability and efficiency (Niazkar et al., 2024). It builds additive decision trees sequentially, with each tree trained to minimize the residual error from the previous iteration (Li, 2022). The objective function combines a differentiable convex loss function L and a regularization term Ω to prevent overfitting (Ekanayake et al., 2022):

$$L\ (\phi) = \sum_{i=1}^{n} l\left(y_i, \hat{y}_i^{(t)}\right) + \sum_{k=1}^{t} \Omega\left(f_k\right) \tag{1}$$

$$\Omega\ (f) = \gamma T + \frac{1}{2}\lambda \|w\|^2 \tag{2}$$

where L is the loss function, y_i and $\hat{y}_i$ are observed and predicted SM values, f_k is the kth regression tree, T is the number of leaves, w is the vector of leaf weights, and γ and λ are regularization parameters. The model updates predictions iteratively using the gradient of the loss function:

$$\hat{y}_i^{(t)} = \hat{y}_i^{(t-1)} + \eta f_t(x_i) \tag{3}$$

where η is the learning rate and $f_t(x_i)$ is the prediction of the new tree.

2. LGBM: LGBM builds upon gradient boosting but introduces two innovations (Kim et al., 2025): (i) a leaf-wise growth strategy with depth constraints to reduce loss more efficiently and (ii) histogram-based feature binning to accelerate training on large datasets. The optimization objective is similar to XGBoost, minimizing loss function L, but with efficient approximation methods that enable faster convergence while maintaining accuracy. The leaf-wise strategy tends to produce deeper trees with better fit to complex patterns, making it suitable for capturing nonlinear interactions between SM and weather variables.
3. Random forest (RF): RF is a bagging-based ensemble model that builds multiple decision trees on bootstrapped samples of the training data and averages their predictions to reduce variance and improve generalization. It is effective at mitigating overfitting and provides stable results across different cover crop treatments (Breiman, 2001). Therefore, it is employed as a baseline model.

3.2 | DL models

DL models were employed to complement the tree-based ML approaches by leveraging their ability to capture temporal dependencies and nonlinear interactions in sequential SM, weather, and water input data. DL architectures can automatically learn latent representations from raw time series, making them particularly suitable for high-frequency SM prediction (Bakhshian et al., 2025). These models are also more capable of modeling complex temporal dynamics introduced by weather variability and sudden water inputs (Ley et al., 2024).

1. Multi-layer perceptron (MLP): An MLP is a fully connected feedforward neural network where each layer applies a nonlinear transformation to its inputs (Ahmadnejad et al., 2024). Given an input vector $\mathbf{x}$, the hidden layers compute:

$$h^{(l)} = \sigma W^{(l)} h^{(l-1)} b^{(l)} \tag{4}$$

where $W^{(l)}$ and $b^{(l)}$ are trainable weights and biases, and σ is an activation function (Pinkus, 1999). MLPs serve as a baseline to benchmark the benefits of temporal modeling of all features.

2. Temporal convolutional network (TCN): TCNs use causal and dilated convolutions to capture long-range dependencies in sequences (Yang et al., 2023). For an input sequence x_t, the dilated convolution at layer l with filter size k and dilation factor d is defined as:

$$h_t^{(l)} = \sum_{i=0}^{k-1} w_i^{(l)} x_{t-d\cdot i} \tag{5}$$

where $h_t{}^{(l)}$ is the hidden activation, $w_i{}^{(l)}$ is the convolution filter weights, and $x_{t\text{-}d\bullet i}$ is the input at lag $d\bullet i$. This structure ensures that predictions at time t only depend on past obser-

vations (Hewage et al., 2020). While the dilated convolutions allow TCNs to efficiently capture SM dynamics spanning over multiple temporal scales without requiring very deep (e.g., many layered) architectures.

3. LSTM: LSTMs are recurrent neural networks designed to capture sequential dependencies while mitigating vanishing gradient problems (Yin et al., 2023). At time step t, the LSTM updates its cell state c_t and hidden state h_t using:

$$f_t = \sigma\left(W_f\left[h_{t-1}, x_t\right] + b_f\right) \tag{6}$$

$$i_t = \sigma\left(W_i\left[h_{t-1}, x_t\right] + b_i\right) \tag{7}$$

$$o_t = \sigma\left(W_o\left[h_{t-1}, x_t\right] + b_o\right) \tag{8}$$

$$o_t = \sigma\left(W_o\left[h_{t-1}, x_t\right] + b_o\right)\ c_t = f_t \odot c_{t-1} + i_t \odot \tanh\left(w_c\left[h_{t-1}, x_t\right] + b_c\right) \tag{9}$$

$$h_t = o_t \odot \tanh\left(c_t\right) \tag{10}$$

where f_t, i_t, and o_t represent the forget, input, and output gates, respectively (Gers et al., 2000). LSTMs are particularly effective in modeling the persistence and lag effects of SM.

4. Transformer: Transformers rely on self-attention mechanisms rather than recurrence to capture dependencies across all time steps simultaneously (Yin et al., 2024; Yin, Carroll, et al., 2025; Yin, Ikard, et al., 2025). For input embeddings $X \in \mathbb{R}^{n\times d}$, the scaled dot-product attention is computed as:

$$\text{Attention}\ (Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \tag{11}$$

where $Q = XW_Q$, $K = XW_K$, and $V = XW_V$ are query, key, and value projections. The self-attention mechanism allows Transformers to learn flexible temporal relationships, making them suitable for capturing both short- and long-term drivers of SM variability (Han et al., 2021).

3.3 | SHAP analysis

To interpret model predictions and quantify the contribution of input variables, we employed SHAP (Lundberg & Lee, 2017), an XAI framework rooted in cooperative game theory. SHAP is particularly suitable in this study because it provides both local explanations (how each feature influences an individual prediction) and global explanations (how features behave on average across the dataset) (Mitchell et al., 2022). Furthermore, SHAP integrates well with tree-based models through the TreeSHAP algorithm (Lundberg et al., 2020), which computes exact Shapley values efficiently for ensemble methods such as XGBoost and LGBM.

The foundation of SHAP is the Shapley value, which attributes a model's prediction to each input feature by considering all possible feature coalitions. For a model f with feature set F, the Shapley value for feature i is defined as (Lundberg & Lee, 2017):

$$\Phi_i\ (f, x) = \sum_{s\subseteq F\setminus\{i\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!}\left[f_{S\cup\{i\}}(x_{S\cup\{i\}}) - f_s\left(x_s\right)\right] \tag{12}$$

where S is a subset of features excluding i, $f_S(x_S)$ is the model prediction using only the feature in S, and $\Phi_i(f, x)$ is the marginal contribution of feature I to the prediction, for instance, x.

In practice, exact computation is intractable for large feature sets, but SHAP uses efficient approximations and model-specific algorithms (e.g., TreeSHAP for tree ensembles). In this study, SHAP was applied in three complementary ways:

1. Feature aggregation: SHAP values were aggregated across the dataset (Bhatt et al., 2020) to rank the relative importance of meteorological and water input variables, including precipitation, irrigation, temperature, humidity, wind speed, and SR.
2. Time-step aggregation: For sequential models, SHAP values were further aggregated across lagged time steps to identify the most critical temporal ranges for prediction (Sood & Craven, 2022).
3. Visualization: To enhance interpretability, several SHAP-based plots were generated, including Beeswarm plots that summarize the distribution of feature contributions across all predictions; timestep-feature heatmaps that reveal how feature importance evolved dynamically across time lags; dependency plots that capture nonlinear feature effects and interactions like between water input and SR (Scheda & Diciotti, 2022).

Together, these analyses provided interpretable insights into both the predictive mechanics of the models and the external features governing SM variability across daily and high-frequency timescales. Importantly, the reliability of

SHAP interpretations depends on the predictive accuracy of the underlying model (Wang et al., 2022). When the model captures SM dynamics with high fidelity, the resulting SHAP values more faithfully reflect true feature contributions and interactions, thereby offering trustworthy insights into the processes driving variability. Conversely, if predictive performance is poor, SHAP explanations may be less meaningful, as they would be based on spurious model patterns rather than real environmental signals. In this study, SHAP analyses were therefore interpreted in conjunction with model performance metrics, ensuring that feature attributions were grounded in well-performing predictive models.

3.4 | Training and evaluation configuration

All models were trained and evaluated under consistent experimental protocols. The limited training size for data-driven modeling raises the potential for overfitting. To mitigate this, the dataset was split into training, validation, and testing subsets with an 8:1:1 ratio for both daily and 5-min data. Model training was monitored on the validation set, and early stopping was employed to prevent overfitting. In addition, performance consistency between training and test sets was verified through cross-validation.

1. ML models: Tree-based models were optimized with root mean squared error (RMSE) as the objective function. XGBoost and LGBM used gradient boosting with shrinkage ($\eta = 0.05$) and subsampling to enhance generalization, while RF relied on bootstrap aggregation. Hyperparameters were tuned to typical values for SM forecasting, with large estimator caps and early stopping to account for dataset size.
2. DL models: Neural networks (MLP, TCN, LSTM, and Transformer) were trained with mean squared error (MSE) loss and optimized with Adam (learning rate = 1×10^{-3}). A mini-batch size of 128 and a maximum of 50 epochs were used, with early stopping after 10 epochs of non-improvement. Architectures varied in depth and capacity: MLP employed three dense layers, TCN stacked dilated convolutions with residual connections, RNNs (GRU/LSTM) used stacked hidden units, and the Transformer incorporated multi-head self-attention with positional encoding.

The RMSE and MSE used for loss and evaluation are defined as:

$$\mathrm{MSE} = \frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2,\ \mathrm{RMSE} = \frac{1}{n}\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2 \quad (13)$$

And the R^2 used for evaluation is defined as:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}\left(y_i - \hat{y}_i\right)^2}{\sum_{i=1}^{n}\left(y_i - \bar{y}\right)^2} \quad (14)$$

where y_i and $\hat{y}_i$ denote observed and predicted SM, respectively, and $\bar{y}$ is the mean of the observations.

These metrics are suitable in this study because MSE and RMSE penalize larger errors more heavily, which is important for capturing abrupt changes in SM due to rainfall or irrigation, while R^2 quantifies the proportion of variance explained by the model, providing an interpretable measure of predictive fidelity across treatments and temporal scales (Yin et al., 2024). Key model configurations are summarized in Table 2. The models are implemented by PyTorch in a Windows system and trained on an Intel i9-13900K @ 3.00 GHz CPU with 128 GB RAM and an Nvidia RTX 4090 GPU to accelerate computation.

4 | RESULTS

4.1 | Lower frequency analysis

All available features, including SM, ST, Tmean, RHmean, Wmean, WaT, and SR from the past 14 days in the daily dataset, were used to predict the SM of the following day for each cover crop treatment (Figure 4). Training the ML models was computationally efficient, and was completed almost instantaneously. Although the training set contained fewer than 700 daily samples per treatment, which is relatively small compared to typical requirements for tree-based ensemble models (often thousands of samples are recommended for stable generalization in regression tasks). XGBoost and LGBM obtained reasonably satisfactory results, with most RMSE values less than 0.05 $m^3\ m^{-3}$ (Figure 4b,c). The R^2 values are generally close to 0.4, which is not high but acceptable given the stochastic occurrence of rainfall and irrigation events, as well as the inherent noise in SM measurements. On the other hand, RF performance is notably weaker, with most R^2 values around 0.2 and larger RMSE values. This difference likely reflects the sequential learning strategy of boosting: for autoregressive and day-to-day correlated variables such as SM, boosting allows each new tree to iteratively correct the errors of previous trees, better capturing nonlinear temporal dependencies than bagging-based RF.

To assess generalization and potential overfitting, we report train/test R^2 and RMSE for the primary benchmark (i.e., LGBM) in Table S1. While the train–test performance gap (training $R^2 \approx 0.9$ compared to testing $R^2 \approx 0.4$) appears large, this pattern was consistent across all independently

TABLE 2 Key model configurations.

Model	Task	Key parameters
RF	Daily	n_estimators = 600, min_samples_leaf = 3
XGBoost	Daily	n_estimators = 500, learning_rate = 0.05, max_depth = 5, subsample = 0.8, colsample_bytree = 0.8
LGBM	Daily	n_estimators = 500, learning_rate = 0.05, num_leaves = 63, min_data_in_leaf = 50, subsample = 0.8, colsample_bytree = 0.8
XGBoost	5 min	n_estimators = 3000, early stop patience = 150, learning_rate = 0.05, max_depth = 6, subsample = 0.8, colsample_bytree = 0.8
LGBM	5 min	n_estimators = 3000, early stop patience = 150, learning_rate = 0.05, num_leaves = 63, min_data_in_leaf = 50, subsample = 0.8, colsample_bytree = 0.8
MLP	5 min	hidden layers [256,128,64], dropout = 0.1
TCN	5 min	dilations = (1…128), filters = 64, kernel size = 3, residual blocks
LSTM	5 min	units = 128, dropout = 0.2, stacked/bidirectional variants
Transformer	5 min	d_model = 128, num_heads = 4, d_ff = 256, depth = 3, dropout = 0.1

Abbreviations: LGBM, light gradient boosting machine; LSTM, long short-term memory; MLP, multi-layer perceptron; RF, random forest; TCN, temporal convolutional network; XGBoost; eXtreme gradient boosting.

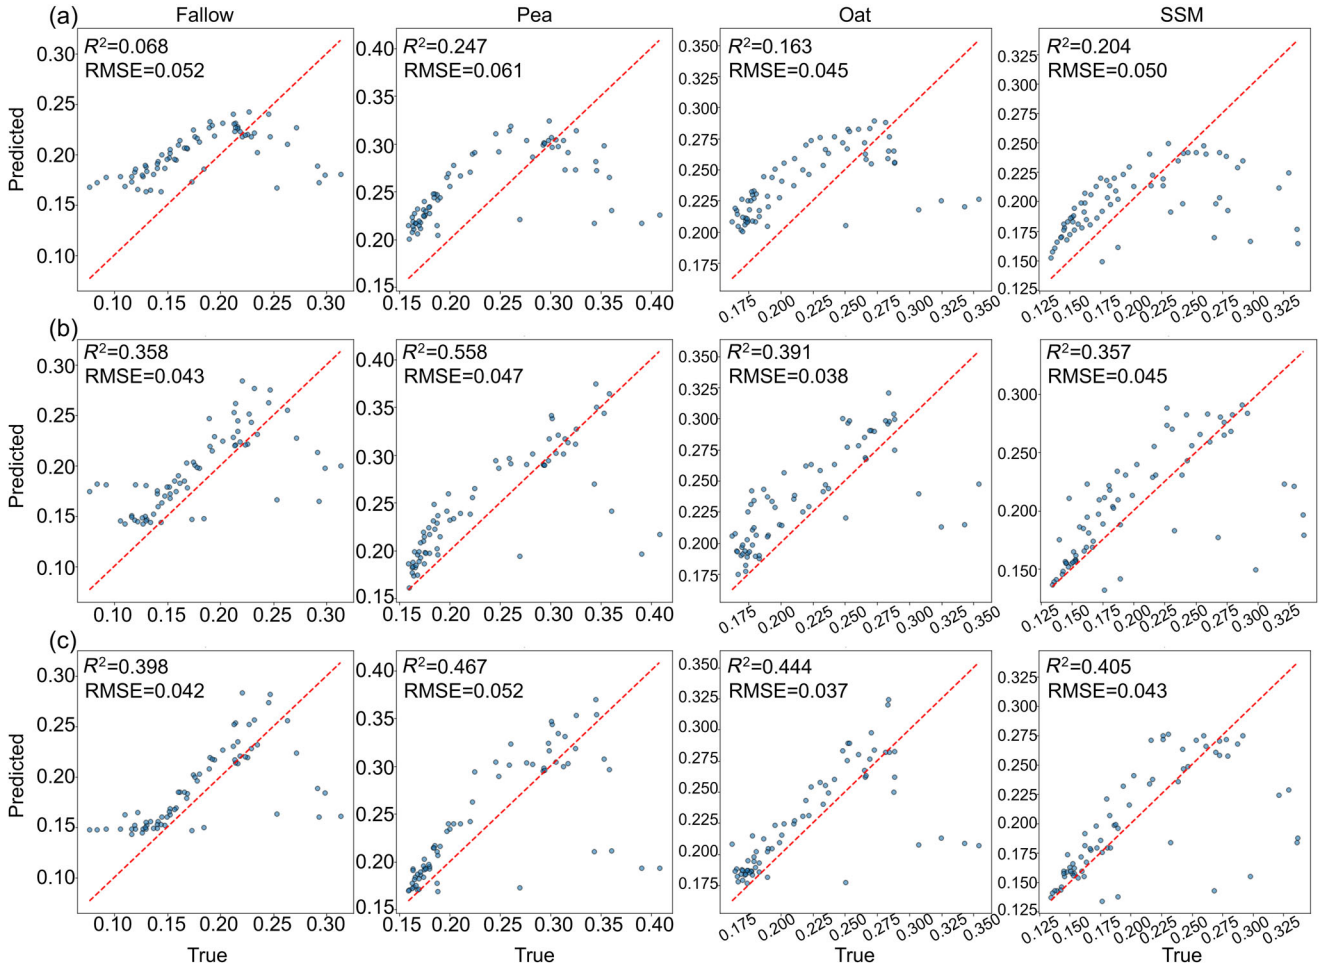

FIGURE 4 Prediction results on the daily testing set for cover crop conditions (each column of plots) using (a) random forest (RF), (b) light gradient boosting machine (LGBM), and (c) eXtreme gradient boosting (XGBoost) (across each row of plots). Each subplot compares observed as true versus predicted soil moisture, with the identity line shown in red. RMSE, root mean squared error; SSM, six-species mixture.

modeled cover crop treatments, indicating that it reflects data characteristics rather than model instability. The gap primarily arises from the stochastic and site-specific variability in rainfall and irrigation events, which strongly and irregularly influence short-term SM changes (Figure 1). Learning-curve inspection (Figure S1) confirmed stable model convergence without major divergence between training and validation losses. Importantly, the consistent test performance across treatments demonstrates that the models retained generalizable predictive capability despite environmental noise. This stability facilitates reliable SHAP-based interpretability and cross-treatment performance diagnostics, providing meaningful insights into feature influence and SM dynamics even under stochastic forcing conditions.

The performance differences among treatments are also evident. SM in the pea plot is the most predictable, achieving an R^2 value of 0.558, while the fallow and SSM plots are less predictable, with R^2 values only up to approximately 0.4. Under identical weather and irrigation conditions, this discrepancy is likely due to differences in canopy cover, root structure, and evapotranspiration dynamics: legumes such as pea tend to moderate soil water fluctuations, whereas fallow plots expose bare soil to rapid evaporative losses and the species mixture (i.e., SSM) introduces greater heterogeneity in root length, density, and distribution, making SM harder to predict. These results justify the selection of XGBoost and LGBM as benchmark models for subsequent high-frequency prediction tasks and as the basis for SHAP analysis to interpret feature and timestep contributions in both daily and short-term prediction settings. SHAP values were aggregated over timesteps and features to provide both temporal and variable-level attributions, which were further visualized using beeswarm and heatmap plots (Figure 5). This analysis allows us to disentangle which drivers and temporal windows exert the strongest influence on SM prediction across different cover crop treatments.

Figure 5a demonstrates that only the most recent 3 days significantly impact predictions across all treatments, with SHAP importance rapidly decaying beyond this short temporal window to below 0.001. This indicates that SM dynamics are strongly autoregressive and dominated by immediate antecedent conditions. Figure 5b shows that, across the 2-week input window, historical SM exerts the greatest influence on predictions, with SHAP values exceeding 0.006. The beeswarm distribution shows that higher SM values (reddish points) are predominantly associated with positive SHAP values, while lower SM values (bluish points) are associated with negative SHAP values. Moreover, the data distribution is wider along the x-axis than for other features, indicating stronger and more variable impacts of SM on the prediction output, thereby confirming its persistence and memory effects. For WaT, SHAP values indicate it is the second-most-influential driver after SM, with contributions concentrated in the most recent few days. The beeswarm plot shows that higher WaT values (reddish points) cluster on the right side of the vertical axis with the maximal value close to 0.002, corresponding to positive contributions that increase predicted SM, while lower WaT values (bluish points) fall slightly on the left side, indicating lower water input quantities (i.e., low irrigation rates) would not be adequate to make positive contributions over longer term under most conditions. The narrower spread compared to SM suggests that WaT has a more direct, yet less variable, impact, consistent with its role as an external water addition. In contrast, the meteorological variables (i.e., Tmean, RHmean, Wmean, and SR) exhibit comparatively smaller SHAP values of ˂0.001, reflecting weaker predictive influence. Their beeswarm distributions are centered closer to zero, with scattered reddish and bluish points indicating occasional positive or negative contributions. This pattern suggests that while weather factors modulate SM, their effects are secondary (relative to antecedent SM and direct water inputs) and are likely mediated by evapotranspiration processes over longer timescales or in specific cover crop contexts.

Figure 5c further clarifies the joint feature–timestep effects. Water input during the most recent 3 days emerges as universally important for all cover crops with distinctively high SHAP values, aligning with the sharp response of SM to irrigation and rainfall events. However, the influences of other variables varies across treatments (e.g., fallow and SSM plots), that appear more sensitive to weather conditions with more variability across weather variables, are likely due to the lack of canopy cover in fallow and the heterogeneous rooting patterns in SSM. In contrast, pea and oat plots show greater stability with very small SHAP values across features other than WaT, which suggests that SM dynamics are more tightly regulated by their own antecedent moisture and water inputs.

Together, these results highlight that daily SM prediction is predominantly associated with short-term memory and direct water additions, while crop management modulates sensitivity to meteorological drivers. This provides an explanation of sensitivities for the performance differences (Figure 4) and further validates the use of SHAP as a tool for linking predictive modeling with process-level understanding. These insights also informed the design of the short-term, high-frequency prediction task, with the past 72 h selected as the lookback window. By leveraging data with finer temporal resolution and greater variability, the subsequent 5-min analysis enabled both improved short-term prediction and a deeper investigation into how feature influences evolution under rapidly changing conditions.

4.2 | Higher frequency analysis

Building on the findings from the daily prediction task, we extended the analysis to the 5-min high-frequency dataset to capture short-term SM dynamics with greater temporal reso-

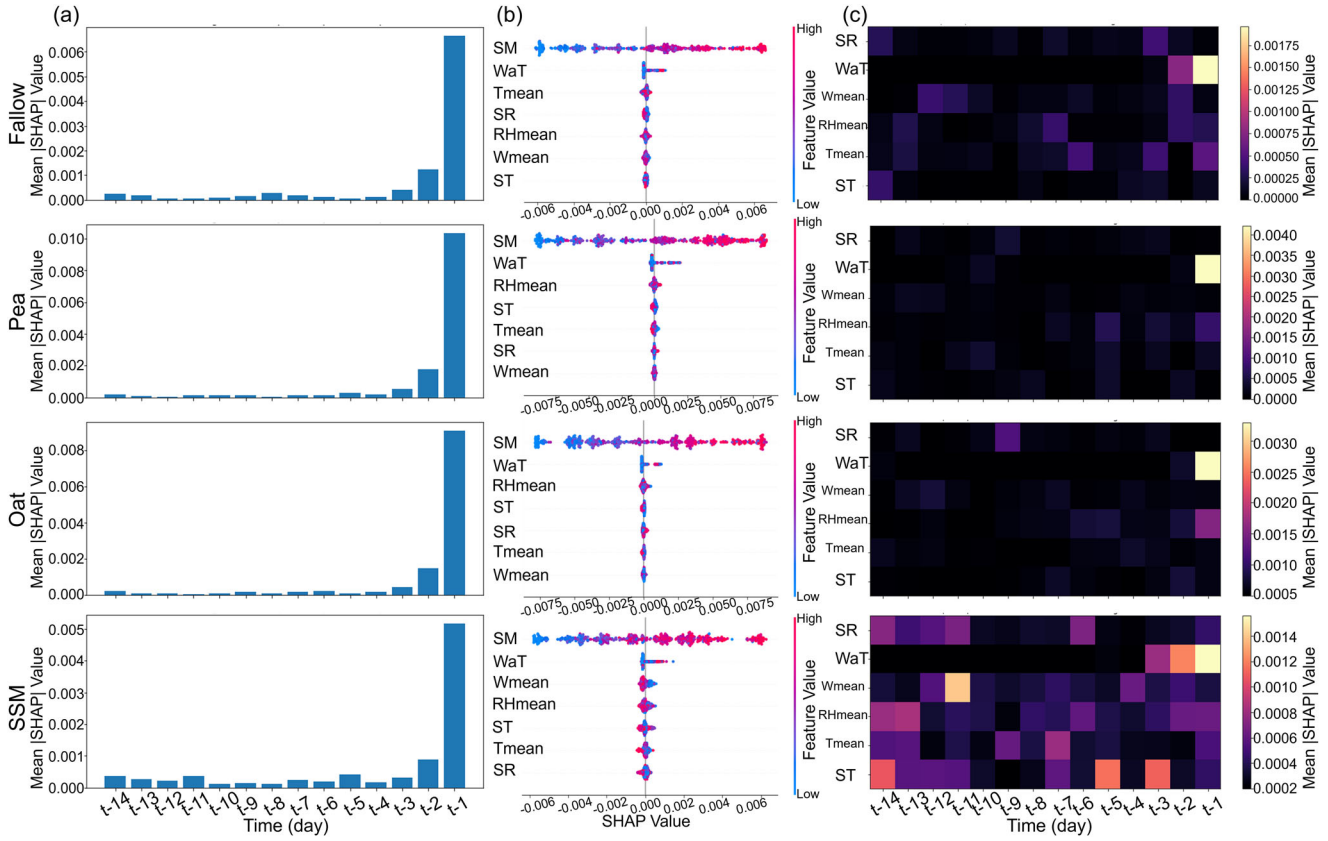

FIGURE 5 SHapley Additive exPlanations (SHAP)-based interpretability analysis of daily soil moisture predictions using light gradient boosting machine (LGBM). Each column of plots includes (a) aggregated SHAP values over timesteps, (b) aggregated SHAP values over features (ranked over mean SHAP values from top to bottom), and (c) feature–timestep heatmaps, shown for each cover crop treatment (each row of plots). SM, soil moisture; SR, solar radiation; SSM, six-species mixture; ST, soil temperature.

lution. In this setting, models were trained to predict SM for the next 8 h using information from the past 72 h, a look-back window heuristically selected based on SHAP-derived timestep importance from the daily analysis (from Figure 5). Compared to daily data, the 5-min dataset exhibits greater variability and more abrupt fluctuations driven by rainfall and irrigation events, making short-term predictions both more challenging and more informative. This higher resolution analysis also allows us to assess not only predictive accuracy with more advanced neural networks but also the temporal alignment of model outputs with observed SM changes. The higher-resolution approach also allows us to further investigate feature influences under rapidly changing conditions and uncover patterns that could not be reliably inferred from the sparser daily data. To this end, both ML (LGBM and XGBoost) and DL models (MLP, TCN, LSTM, and Transformer) were evaluated, with LGBM and XGBoost serving as benchmark methods against which the performance of neural architecture was compared (Figure 6). The results are presented for the Fallow plot, which is the most challenging to predict, as shown previously in Figure 4. Among all models, LGBM demonstrates a clear advantage, producing predictions that are more closely aligned with the identity line and achieving the highest R^2 of 0.671. Both LGBM and XGBoost achieve significant performance improvements compared to using daily data. Transformer and LSTM models perform comparably well, with R^2 values approaching 0.5, whereas MLP and TCN yield unsatisfactory performance with R^2 close to 0.3. This poorer performance likely reflects their limited ability to capture long-term dependencies and nonlinear responses to abrupt water input events.

Interestingly, the differences in RMSE values are relatively minor among the stronger DL models (LGBM, XGBoost, Transformer, and LSTM are all close to 0.045), suggesting that while they achieve similar average error magnitudes, their ability to explain variance (as captured by R^2) diverges more substantially (Figure 6). This pattern suggests that variance in SM fluctuations is better captured by LGBM, even though overall error magnitudes are not substantially different. In terms of computational efficiency, the tree-based models (LGBM and XGBoost) show a clear advantage, completing training in under 5 min. In contrast, DL models are more computationally demanding. LSTM and Transformer required 44 s and 295 s per epoch, on average, respectively, leading to total

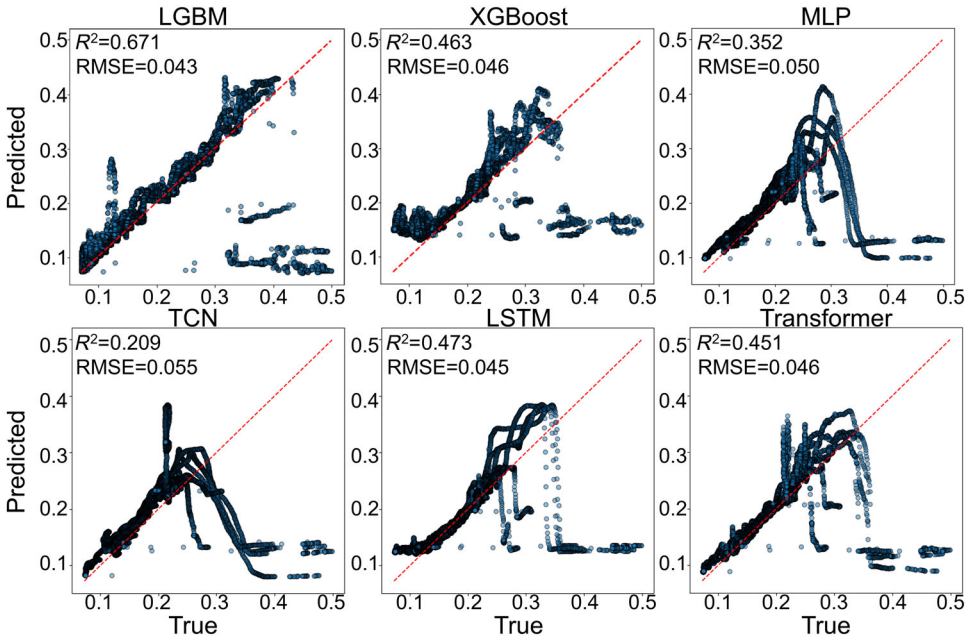

FIGURE 6 Prediction results on the 5-min testing set for the fallow plot using light gradient boosting machine (LGBM), eXtreme gradient boosting (XGBoost), multi-layer perceptron (MLP), temporal convolutional network (TCN), long short-term memory (LSTM), and Transformer. Each subplot shows predicted versus observed soil moisture with the identity line in red, annotated with R^2 and root mean squared error (RMSE) values. Noticeable deviations (across all six models) from the identity line when soil moisture (SM) is >0.3 were generally caused by irrigation or few heavy rainfalls, indicating the inability to capture soil moisture variation patterns under large water input due to scarce data that include such events.

training times exceeding 30 min and 4 h. This highlights the trade-off between accuracy, interpretability, and efficiency when selecting models for operational use. The 72-h lookback choice was informed by, but not independently validated against, the daily SHAP results. To evaluate the robustness of this assumption, sensitivity testing of the lookback length (48, 72, and 96 h) produced negligible improvements in model accuracy ($\Delta R^2 < 0.01$ with the increase of lookback length) on the LGBM benchmark (Table S2), indicating that temporal patterns identified are robust to moderate variations in input window length. The 72-h configuration thus represents a practical balance, providing satisfactory predictive performance while avoiding unnecessary computational overhead. Additionally, across all models, the largest discrepancies between predictions and observations appear as data clusters below the identity line. These represent SM underestimations associated with abrupt fluctuations caused by external water inputs (irrigation or rainfall). Such events remain particularly challenging to capture accurately due to their stochastic nature.

To further evaluate the temporal prediction performance of promising models, we then compare LGBM (daily and 5 min) and Transformer results directly in the time domain in Figure 7. The 5-min predictions (Figure 7b,c) align more closely with the observed SM curve compared to the daily model (Figure 7a), where the prediction curve shows a noticeable rightward shift relative to the observations. This demonstrates that higher temporal resolution not only improves accuracy but also reduces the lag effect, thereby increasing the timeliness of the predictions. On the other hand, all major predicted SM peaks also align temporally with recorded irrigation events listed in Table 1, further confirming that the models accurately capture short-term responses to water inputs, even though this correspondence should be interpreted as temporal alignment rather than physical causation. Between the two high-frequency models, the Transformer (Figure 7c) shows the least lag, with predictions nearly synchronous with observed fluctuations. However, responsiveness to extreme spikes caused by water input events is limited, as peak values approaching 0.5 in the observations are underestimated below 0.3. In contrast, the LGBM 5-min model (Figure 7b) achieves a more balanced performance. It better reproduces the magnitude of peak responses (~0.4), while still maintaining close overall alignment, albeit with slightly more lag than the Transformer. These results highlight a trade-off between timeliness (Transformer) and event responsiveness (LGBM), suggesting that ensemble tree-based and attention-based neural architectures capture complementary aspects of short-term SM dynamics. SHAP analysis was conducted on the 5-min LGBM model results to further evaluate feature influence over the 3-day window.

Figure 8a shows that SM itself remains the most influential feature for predicting moisture content, with SHAP values reaching the largest magnitudes (up to 0.0002 in model

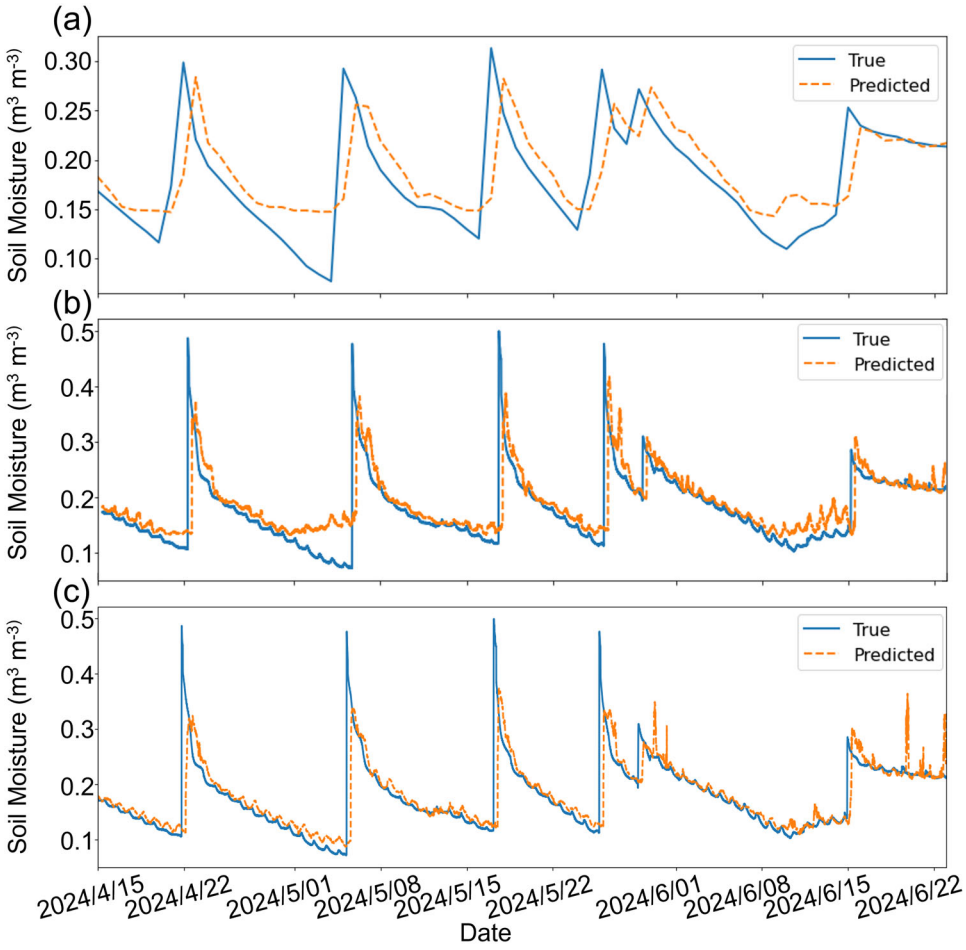

FIGURE 7 Temporal prediction results for the Fallow plot using (a) light gradient boosting machine (LGBM) daily, (b) LGBM 5-min, and (c) Transformer 5-min models. Blue solid lines indicate observed soil moisture, and orange dashed lines indicate model predictions.

output units), consistent with the persistence and memory effects seen in the daily analysis. Because SHAP values are computed continuously across the time series, their interpretation reflects the overall influence of recent water inputs rather than individual irrigation events. Therefore, the water input variable (WaT), which aggregates rainfall and irrigation into a continuous feature, ranked lower in aggregate but exerts a consistently strong positive influence, as indicated by the rightward skew of SHAP values. Unlike the daily scale, SR emerges as the most aggregated feature contributor across the 72-h window in most treatments, despite its narrower SHAP range. Variability across treatments is also more pronounced than in the daily case: for example, air temperature (Tmean) shows a stronger negative effect in the fallow plot, suggesting enhanced evaporation driven by temperature in the topsoil. In contrast, humidity (RHmean) shows a stronger positive effect in the SSM plot, implying that SM is closely related to evapotranspiration, as higher RHmean helps mitigate crop transpiration. Together, these findings indicate that over shorter 3-day periods, SM dynamics are not only dominated by autoregression and water additions but also become more sensitive to radiative and atmospheric drivers.

Recent WaT values remain dominant predictors, but SR also displays recurrent bands every 24 h, with strong influence across the entire 72-h horizon (Figure 8b). Additionally, both air and STs exhibit localized peaks in the most recent hours, suggesting their role in modulating rapid evaporative losses immediately following water input events. This implies that short-term predictions depend strongly on immediate water inputs and are shaped by near-real-time atmospheric forcing. Specifically, WaT exhibits two clear peaks in positive influence on SM predictions, centered around 4–5 cm/day and 7–9 cm/day, with SHAP values approaching 1.0×10^{-5} (Figure 8c). These thresholds likely correspond to irrigation or rainfall regimes that have the most pronounced impacts on near-surface moisture conditions. The first peak (4–5 cm/day) may represent typical irrigation or moderate rainfall events that quickly recharge the upper soil layer and are well captured by the sensors. The second peak (7–9 cm/day) likely reflects irrigation or larger rainfall that exceeds infiltration capacity, leading to transient soil saturation, surface runoff,

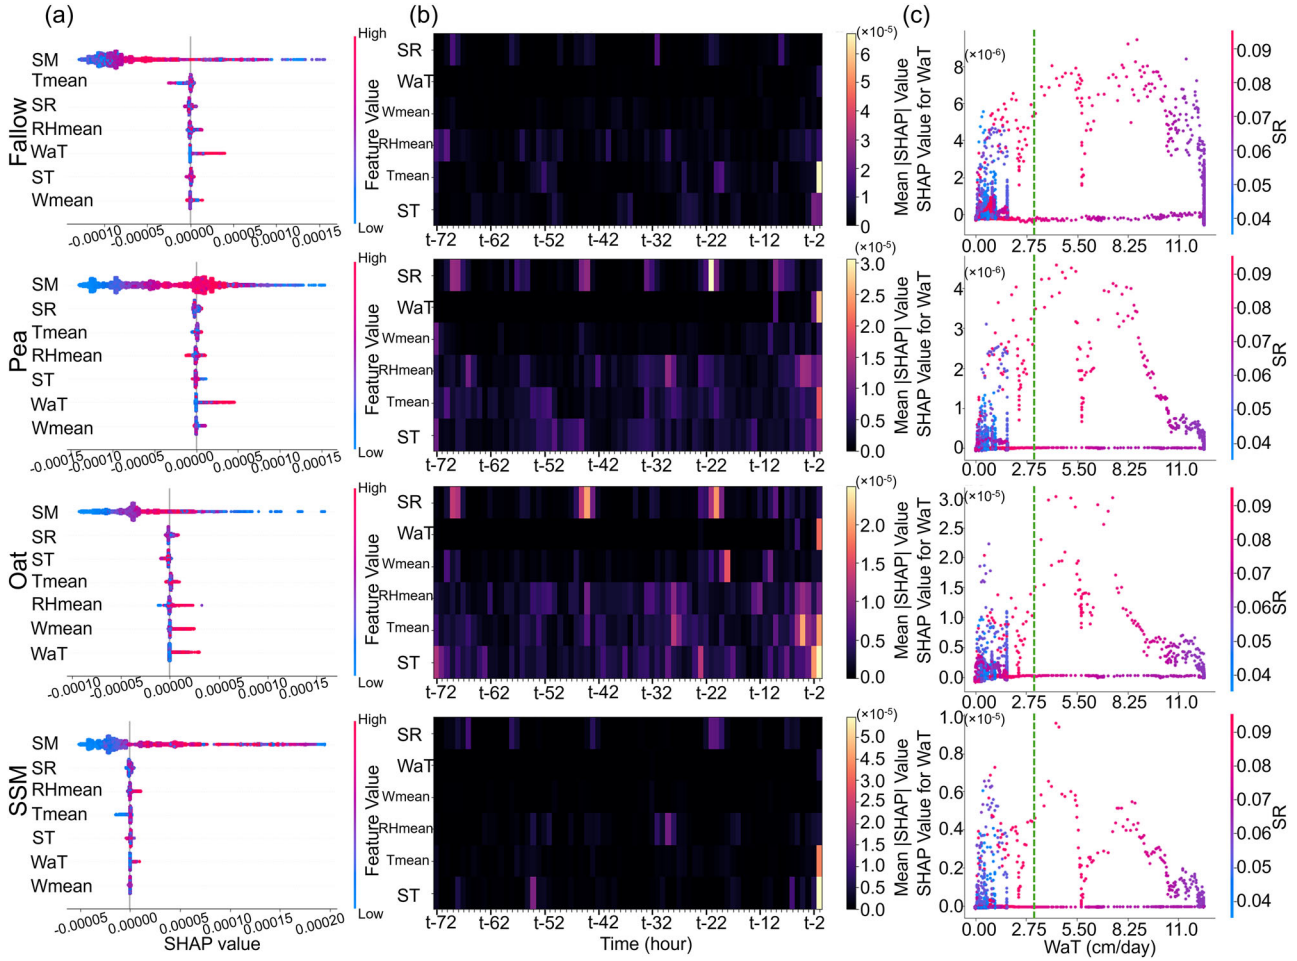

FIGURE 8 SHapley Additive exPlanations (SHAP) analysis of 5-min prediction results using the light gradient boosting machine (LGBM) model for different cover crop treatments (each row of plots). Beeswarm plots (a) of feature importance aggregated over the 72-h input window are ranked by high to low feature importance, from top to bottom in each plot (left column). Middle column contains feature–timestep heatmaps (b), and right column contains dependency plots (c) of water input (WaT, converted from cm/5 min to cm/day) and solar radiation (SR). The green dashed line represents the average irrigation amount per event. SM, soil moisture; SSM, six-species moisture; ST, soil temperature.

or rapid drainage. Such nonlinear responses create stronger contrasts in SM dynamics, which are then emphasized in the SHAP attribution. By contrast, SR shows only minor modulation of SHAP values across different WaT levels, suggesting that its effect operates more indirectly through evapotranspiration processes. In addition, we found higher SHAP variability for cover crops, including oat (maximal SHAP value $\sim 3.0 \times 10^{-5}$) and SSM (maximal SHAP value $\sim 1.0 \times 10^{-5}$), which implies increased sensitivity in those treatments to water input, improved water infiltration and uptake within 72 h of rainfall/irrigation, and less water loss due to runoff and evaporation. Together, these results highlight that within short 3-day windows, SM variability is primarily driven by a balance between immediate water additions and atmospheric influences, with crop cover influencing sensitivity to secondary weather variables, which could be used to optimize water use by improving water infiltration and uptake.

While SHAP provides valuable interpretability for understanding feature influence, it is essential to note that these attributions represent statistical sensitivities within the model space rather than direct causal mechanisms of soil–plant–atmosphere interactions. Even though temporal autocorrelation is inherent in SM and water input variables, the relatively low cross-feature correlations among weather predictors (Figure 2b,c) suggest that the derived feature influences largely reflect distinct, data-driven relationships within the model space. Accordingly, SHAP interpretations in this study should be viewed as data-driven associations that enhance understanding of feature influence patterns rather than mechanistic controls on soil water processes.

5 | DISCUSSION

5.1 | Method strength and limitations

The results from both daily and 5-min analyses highlight several important aspects of SM dynamics in the semi-arid regions, such as the study area in Clovis, NM. The dominance of recent SM and water inputs in the SHAP analyses indicates strong short-term memory of near-surface soils and their immediate responsiveness to irrigation and rainfall pulses. This finding is consistent with previous research in water-limited systems, where SM is primarily controlled by antecedent water status and discrete recharge events, rather than gradual meteorological forcing alone (Teshome et al., 2024). At the same time, the emergence of SR and temperature as important factors in the high-frequency analysis underscores the role of evaporative demand in driving rapid drying cycles following wetting events.

The results for this new dual-resolution modeling framework show performance that aligns with that of Teshome et al. (2024), who reported that gradient-boosting algorithms such as XGBoost and LGBM consistently outperformed other ML and DL models in SM prediction tasks. However, while their analysis was limited to a single sub-hourly scale, our dual-resolution framework extends this approach by jointly evaluating daily and high-frequency dynamics and incorporating SHAP-based interpretability to quantify feature influence across temporal scales. Advantages of ensemble boosting (Alahmad et al., 2025) and SHAP analysis (Wang et al., 2024) for hydrological prediction have also been reported in other regions and climates, suggesting that these models provide a robust foundation for soil–water forecasting under both humid and semi-arid conditions (Rahmati et al., 2024). Prior studies have also shown that multiscale hydrologic modeling may include multiple length scales, multiple time scales, or both (Scheibe et al., 2015). The results of the new dual-resolution approach developed here advance this understanding by demonstrating that predictive sensitivity to meteorological and management variables is time scale-dependent, with SR and short-term water inputs exerting greater influence on finer temporal resolutions in the site studied.

From an agronomic perspective, the ability of these models to capture both the persistence of soil water and the timing of sharp fluctuations is highly relevant for irrigation management (Matei et al., 2017). The predictive framework for ML/DL models incorporating SHAP analysis demonstrated here better captures and predicts low moisture levels than high moisture conditions, providing reliable, timely guidance to farmers on when to irrigate. It can also help identify periods of heightened water sensitivity, enabling more efficient irrigation scheduling to reduce losses from runoff or deep percolation. Moreover, the observed differences among cover crop treatments, with pea plots showing greater stability in water storage and fallow or SSM plots exhibiting higher variability, suggest that cover cropping strategies can influence not only soil water retention but also the predictability of SM itself under various meteorological conditions (Florence & McGuire, 2020). For example, farmers adopting mixed cover cropping may need to consider the water-use characteristics of the constituent species when making water management decisions. Such insights may help farmers balance crop water use, soil health, and management practices under semi-arid conditions. In addition, the framework provides actionable insights into irrigation management and cover crop selection in the semi-arid site examined here. The ability to predict SM with short lead times, identify critical water input thresholds (4–5 and 7–9 cm/day), and quantify treatment-specific sensitivities can help optimize irrigation scheduling, reduce water losses, and improve crop water-use efficiency. The observed differences among treatments further highlight the role of cover crops in modulating SM predictability under semi-arid conditions.

The methodological design presented here also offers notable strengths. By combining ML and DL models with SHAP-based interpretability and a dual-temporal resolution strategy, the study provides both predictive skill and process-level understanding. Tree-based boosting models (LGBM and XGBoost) proved especially effective at capturing nonlinear patterns and providing efficient, interpretable predictions of moisture content, while DL models added value by reducing lag in high-frequency predictions and highlighting complex feature interactions. SHAP analysis further contributed by disentangling the relative roles of soil, weather, and management factors, indicating consistent short-term memory effects of SM and identifying treatment-specific sensitivities that would otherwise remain hidden in black-box model outputs.

Although SHAP analysis quantifies model sensitivities and data-driven associations rather than mechanistic fluxes of infiltration, evapotranspiration, or plant uptake, rather than physical causality. The resulting feature attributions remain valuable for both scientific interpretation and practical management. The analysis identifies which variables most strongly affect SM predictability, informs irrigation scheduling, and highlights periods of heightened water sensitivity across crop treatments. The identified sensitivities, such as the dominance of recent water input and SR, are actionable even without formal causal validation. Therefore, SHAP serves as an effective bridge between data-driven prediction and agronomic interpretation, providing a transparent, quantitative basis for evaluating feature influence and improving decision support in semi-arid agriculture.

On the other hand, several limitations should be acknowledged. First, predictive accuracy was constrained by dataset size, particularly at the daily scale, which limited the ability of more data-intensive DL architectures to realize their

full potential. Second, while the models captured dominant drivers and short-term memory effects, they may not fully account for deeper soil processes such as percolation beyond the sensor depth or plant water uptake at deeper rooting zones. Third, high-frequency fluctuations associated with irrigation or large rainfall events remained difficult to capture due to the inherent stochasticity and sparseness of these inputs, as rainfall and irrigation are common but scarce in the studied semi-arid region. Finally, although SHAP provided valuable interpretability, its reliability in feature attributions depends on the quality of the underlying models, and caution is warranted when interpreting results from models with weaker performance. While model performance and SHAP-based insights are consistent across treatments, future research with multi-site datasets and deeper soil profiles is needed to confirm model robustness and extend generalization.

5.2 | Relation to physical water-balance models

The statistical relationships identified by this proposed ML + DL + SHAP framework can be interpreted in the context of established hydrologic processes represented in physically based models such as HYDRUS-1D and SWAT. In similar semi-arid soils, HYDRUS-1D simulations have shown that near-surface water content responds sharply to irrigation or rainfall pulses followed by rapid drainage and evaporation under high SR (Er-Raki et al., 2021). These infiltration–evaporation cycles described in process models align with the short-term SHAP patterns identified in this study, in which the most recent water-input events and strong radiative forcing govern soil-moisture variations at sub-daily scales. Similarly, SWAT simulations for the southern High Plains indicate that evapotranspiration, canopy shading, and surface storage dominate soil-water variability under limited precipitation, with slow percolation or lateral flow (Chen et al., 2018). It should be noted, however, that several of the direct measurements of individual water-balance components (e.g., infiltration rate, percolation, evapotranspiration, or canopy interception) were not collected in this study, and therefore, mechanistic inference remains limited. The SHAP-based interpretations should thus be viewed as data-driven sensitivities within the model space, complementary to rather than substitutive of process-based modeling (Dai et al., 2025; Xi et al., 2025).

6 | CONCLUSIONS

This study developed a dual-resolution framework integrating ML, DL, and SHAP-based interpretability to predict and explain SM dynamics under various cover crop treatments in a semi-arid environment. Tree-based boosting models (LGBM and XGBoost) achieved the highest predictive accuracy and computational efficiency, while DL models, especially Transformer and LSTM, effectively reduced temporal lag in high-frequency predictions. SHAP analysis provided data-driven insights into feature sensitivities across timescales, indicating that recent SM and water input dominate daily variability, whereas SR and temperature exert stronger short-term influences.

These results demonstrate that combining predictive modeling with XAI enhances both accuracy and interpretability, offering a practical basis for data-informed irrigation and soil and water management. However, because analyses are limited to surface (5 cm) moisture at a single site in Clovis, NM, current results should be interpreted as site-specific until verified across additional locations and depths. Looking ahead, the framework can be further enhanced by coupling the data-driven models with physically based hydrological or soil–plant–atmosphere simulations to enable causal evaluation of feature influences identified by SHAP. Such hybrid modeling would allow comparison between model-derived sensitivities and process-based fluxes such as infiltration, evapotranspiration, and root water uptake, thereby improving both physical interpretability, mechanistic validation, and transferability across sites.

AUTHOR CONTRIBUTIONS

Huichao Yin: Conceptualization; data curation; formal analysis; investigation; methodology; software; validation; visualization; writing—original draft. **Prakriti Bista**: Conceptualization; data curation; formal analysis; investigation; methodology; resources; supervision; validation; writing—review and editing. **Rajan Ghimire**: Conceptualization; data curation; investigation; methodology; project administration; resources; writing—review and editing. **Hui Yang**: Validation; writing—review and editing. **Kenneth C. Carroll**: Conceptualization; funding acquisition; methodology; project administration; resources; supervision; validation; writing—review and editing.

ACKNOWLEDGMENTS

This work was supported by the United States Department of Agriculture-National Institute of Food and Agriculture (grant # 2022-67019-36106 and 2024-68017-42789). Additional support by NMSU's Agricultural Experiment Station and Department of Energy Environmental Management's Minority Serving Institution Partnership Program (MSIPP) managed by the Savannah River National Laboratory under BSRA contract TOA 800002114 is appreciated. We appreciate the ASC Clovis field staff's help in maintaining plots and routine farm operations.

CONFLICT OF INTEREST STATEMENT

This authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

Data are made available in http://doi.org/10.6084/m9.figshare.30263209. Method implementation, trained model weight, and analyzing result data are available at http://doi.org/10.6084/m9.figshare.30263518.

ORCID

Huichao Yin https://orcid.org/0000-0001-6172-5580
Prakriti Bista https://orcid.org/0000-0002-3134-9702
Rajan Ghimire https://orcid.org/0000-0002-6962-6066
Hui Yang https://orcid.org/0000-0002-8518-7972
Kenneth C. Carroll https://orcid.org/0000-0003-2097-9589

REFERENCES

Adebayo, O., Singh, A., Bista, P., Angadi, S., & Ghimire, R. (2025). Compost addition improves soil water storage and crop water productivity in cover crop integrated sorghum production system under a limited irrigation management. *Irrigation Science*, *43*, 1559–1573. https://doi.org/10.1007/s00271-025-01011-2

Ahmadnejad, S., Nadi, M., & Aghelpour, P. (2024). Numerical estimation of surface soil moisture by machine learning algorithms in different climatic types. *Pure and Applied Geophysics*, *181*(7), 2149–2175. https://doi.org/10.1007/s00024-024-03508-x

Alahmad, T., Neményi, M., & Nyéki, A. (2025). Soil moisture content prediction using gradient boosting regressor (GBR) model: Soil-specific modeling with five depths. *Applied Sciences*, *15*(11), 5889. https://doi.org/10.3390/app15115889

Bakhshian, S., Zarepakzad, N., Nevermann, H., Hohenegger, C., Or, D., & Shokri, N. (2025). Field-scale soil moisture dynamics predicted by deep learning. *Advances in Water Resources*, *201*, 104976. https://doi.org/10.1016/j.advwatres.2025.104976

Bhatt, U., Weller, A., & Moura, J. M. (2020). Evaluating and aggregating feature-based model explanations. *arXiv*. https://doi.org/10.48550/arXiv.2005.00631

Breiman, L. (2001). Random forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/a:1010933404324

Chen, J. F., Sainju, U. M., Liu, R., Tan, G. Y., Wen, M. M., Zhao, J., Pu, J. L., Feng, J. R., & Wang, J. (2024). Winter legume cover crop with adequate nitrogen fertilization enhance dryland maize yield and water-use efficiency. *Agricultural Water Management*, *306*, 109209. https://doi.org/10.1016/j.agwat.2024.109209

Chen, Y., Marek, G. W., Marek, T. H., Brauer, D. K., & Srinivasan, R. (2018). Improving SWAT auto-irrigation functions for simulating agricultural irrigation management using long-term lysimeter field data. *Environmental Modelling & Software*, *99*, 25–38. https://doi.org/10.1016/j.envsoft.2017.09.013

Cheng, S., Xue, W., Gong, X., Hu, F., Yang, Y., & Liu, M. (2024). Reconciling plant and microbial ecological strategies to elucidate cover crop effects on soil carbon and nitrogen cycling. *Journal of Ecology*, *112*(12), 2901–2916. https://doi.org/10.1111/1365-2745.14431

Cui, X., Xu, G., He, X. F., & Luo, D. Q. (2022). Influences of seasonal soil moisture and temperature on vegetation phenology in the Qilian Mountains. *Remote Sensing*, *14*(15), 3645. https://doi.org/10.3390/rs14153645

Dai, Z., Zhan, C., Yin, H., Chen, J., Xu, L., Xia, Y., Yang, S., Chen, W., Cao, M., Du, Z., Zhang, X., Yan, B., Ma, Y., Wang, H., Moeini, F., Soltanian, M. R., Thanh, H. V., & Carroll, K. C. (2025). Incorporating deep learning into hydrogeological modeling: Advancements, challenges, and future directions. *Journal of Geophysical Research: Machine Learning and Computation*, *2*, e2025JH000703. https://doi.org/10.1029/2025JH000703

Demir, S., & Sahin, E. K. (2023). An investigation of feature selection methods for soil liquefaction prediction based on tree-based ensemble algorithms using AdaBoost, gradient boosting, and XGBoost. *Neural Computing & Applications*, *35*(4), 3173–3190. https://doi.org/10.1007/s00521-022-07856-4

Ekanayake, I. U., Meddage, D. P. P., & Rathnayake, U. (2022). A novel approach to explain the black-box nature of machine learning in compressive strength predictions of concrete using Shapley additive explanations (SHAP). *Case Studies in Construction Materials*, *16*, e01059. https://doi.org/10.1016/j.cscm.2022.e01059

Er-Raki, S., Ezzahar, J., Merlin, O., Amazirh, A., Hssaine, B. A., Kharrou, M. H., Khabba, S., & Chehbouni, A. (2021). Performance of the HYDRUS-1D model for water balance components assessment of irrigated winter wheat under different water managements in semi-arid region of Morocco. *Agricultural Water Management*, *244*, 106546. https://doi.org/10.1016/j.agwat.2020.106546

Florence, A. M., & McGuire, A. M. (2020). Do diverse cover crop mixtures perform better than monocultures? A systematic review. *Agronomy Journal*, *112*(5), 3513–3534. https://doi.org/10.1002/agj2.20340

Furtak, K., & Wolińska, A. (2023). The impact of extreme weather events as a consequence of climate change on the soil moisture and on the quality of the soil environment and agriculture—A review. *Catena*, *231*, 107378. https://doi.org/10.1016/j.catena.2023.107378

Gers, F. A., Schmidhuber, J., & Cummins, F. (2000). Learning to forget: Continual prediction with LSTM. *Neural computation*, *12*(10), 2451–2471. https://doi.org/10.1162/089976600300015015

Ghimire, R., Ghimire, B., Mesbah, A. O., Sainju, U. M., & Idowu, O. J. (2019). Soil health response of cover crops in winter wheat–fallow system. *Agronomy Journal*, *111*(4), 2108–2115. https://doi.org/10.2134/agronj2018.08.0492

Ghimire, R., Thapa, V. R., Acosta-Martinez, V., Schipanski, M., Slaughter, L. C., Fonte, S. J., Shukla, M. K., Bista, P., Angadi, S. V., Mikha, M. M., Adebayo, O., & Noble Strohm, T. (2023). Soil Health assessment and management framework for water-limited environments: Examples from the Great Plains of the USA. *Soil Systems*, *7*(1), 22. https://doi.org/10.3390/soilsystems7010022

Gu, X. Q., Yao, L., & Wu, L. F. (2023). Prediction of water carbon fluxes and emission causes in rice paddies using two tree-based ensemble algorithms. *Sustainability*, *15*(16), 12333. https://doi.org/10.3390/su151612333

Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., & Wang, Y. (2021). Transformer in transformer. *Advances in Neural Information Processing Systems*, *34*, 15908–15919.

Hewage, P., Behera, A., Trovati, M., Pereira, E., Ghahremani, M., Palmieri, F., & Liu, Y. (2020). Temporal convolutional neural (TCN) network for an effective weather forecasting using time-series data from the local weather station. *Soft Computing*, *24*(21), 16453–16482. https://doi.org/10.1007/s00500-020-04954-0

Kim, S., Kim, H., Kwon, Y., & Nguyen, H. H. (2025). A stand-alone framework for predicting spatiotemporal errors in satellite-based soil moisture using tree-based models and deep neural networks. *Geoscience & Remote Sensing*, *62*(1), 2475572. https://doi.org/10.1080/15481603.2025.2475572

Kornelsen, K. C., & Coulibaly, P. (2014). Root-zone soil moisture estimation using data-driven methods. *Water Resources Research*, *50*(4), 2946–2962. https://doi.org/10.1002/2013wr014127

Ley, A., Bormann, H., & Casper, M. (2024). Linking explainable artificial intelligence and soil moisture dynamics in a machine learning streamflow model. *Hydrology Research*, *55*(6), 613–627. https://doi.org/10.2166/nh.2024.003

Li, X. Q., Gong, M. Y., Dong, J. R., Zhou, Z. Y., Han, B., & Yu, H. L. (2024). Forecasting regional in-situ thermal conductivity of soil based on tree-based ensemble learning. *International Communications in Heat and Mass Transfer*, *159*, 107996. https://doi.org/10.1016/j.icheatmasstransfer.2024.107996

Li, Z. Q. (2022). Extracting spatial effects from machine learning model using local interpretation method: An example of SHAP and XGBoost. *Computers Environment and Urban Systems*, *96*, 101845. https://doi.org/10.1016/j.compenvurbsys.2022.101845

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, *2*(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In I. Guyon et al (Eds.) *Advances in Neural Information Processing Systems* (vol. 30). NeurIPS.

Luo, M., Meng, F. H., Sa, C. L., Duan, Y. C., Bao, Y. H., Liu, T., & De Maeyer, P. (2021). Response of vegetation phenology to soil moisture dynamics in the Mongolian Plateau. *Catena*, *206*, 105505. https://doi.org/10.1016/j.catena.2021.105505

Matei, O., Rusu, T., Petrovan, A., & Mihuţ, G. (2017). A data mining system for real time soil moisture prediction. *Procedia Engineering*, *181*, 837–844. https://doi.org/10.1016/j.proeng.2017.02.475

Mitchell, R., Frank, E., & Holmes, G. (2022). GPUTreeShap: Massively parallel exact calculation of SHAP scores for tree ensembles. *PeerJ Computer Science*, *8*, e880. https://doi.org/10.7717/peerj-cs.880

Niazkar, M., Menapace, A., Brentan, B., Piraei, R., Jimenez, D., Dhawan, P., & Righetti, M. (2024). Applications of XGBoost in water resources engineering: A systematic literature review (Dec 2018-May 2023). *Environmental Modelling & Software*, *174*, 105971. https://doi.org/10.1016/j.envsoft.2024.105971

Nikraftar, Z., Parizi, E., Saber, M., Boueshagh, M., Tavakoli, M., Esmaeili Mahmoudabadi, A., Ekradi, M. H., Mbuvha, R., & Hosseini, S. M. (2025). An interpretable machine learning framework for unraveling the dynamics of surface soil moisture drivers. *Remote Sensing*, *17*(14), 2505. https://doi.org/10.3390/rs17142505

Nilahyane, A., Ghimire, R., Sharma Acharya, B., Schipanski, M. E., West, C. P., & Obour, A. K. (2023). Overcoming agricultural sustainability challenges in water-limited environments through soil health and water conservation: Insights from the Ogallala Aquifer Region, USA. *International Journal of Agricultural Sustainability*, *21*(1), 2211484. https://doi.org/10.1080/14735903.2023.2211484

Peng, Y., Wang, L., Jacinthe, P.-A., & Ren, W. (2024). Global synthesis of cover crop impacts on main crop yield. *Field Crops Research*, *310*, 109343. https://doi.org/10.1016/j.fcr.2024.109343

Pinkus, A. (1999). Approximation theory of the MLP model in neural networks. *Acta Numerica*, *8*, 143–195. https://doi.org/10.1017/S0962492900002919

Qin, J. X., Duan, W. L., Zou, S., Chen, Y. N., Huang, W. J., & Rosa, L. (2024). Global energy use and carbon emissions from irrigated agriculture. *Nature Communications*, *15*(1), Article 3084. https://doi.org/10.1038/s41467-024-47383-5

Qiu, T., Shi, Y. U., Peñuelas, J., Liu, J. I., Cui, Q., Sardans, J., Zhou, F., Xia, L., Yan, W., Zhao, S., Peng, S., Jian, J., He, Q., Zhang, W., Huang, M., Tan, W., & Fang, L. (2024). Optimizing cover crop practices as a sustainable solution for global agroecosystem services. *Nature Communications*, *15*(1), Article 10617. https://doi.org/10.1038/s41467-024-54536-z

Rahmati, M., Amelung, W., Brogi, C., Dari, J., Flammini, A., Bogena, H., Brocca, L., Chen, H., Groh, J., Koster, R. D., Mccoll, K. A., Montzka, C., Moradi, S., Rahi, A., Sharghi S, F., & Vereecken, H. (2024). Soil moisture memory: State-of-the-art and the way forward. *Reviews of Geophysics*, *62*, e2023RG000828. https://doi.org/10.1029/2023RG000828

Scheda, R., & Diciotti, S. (2022). Explanations of machine learning models in repeated nested cross-validation: An application in age prediction using brain complexity features. *Applied Sciences*, *12*(13), 6681. https://doi.org/10.3390/app12136681

Scheibe, T. D., Murphy, E. M., Chen, X., Rice, A. K., Carroll, K. C., Palmer, B. J., Tartakovsky, A. M., Battiato, I., & Wood, B. D. (2015). An analysis platform for multiscale hydrogeologic modeling with emphasis on hybrid multiscale methods. *Groundwater*, *53*, 38–56. https://doi.org/10.1111/gwat.12179

Sharma, A., Jain, A., Gupta, P., & Chowdary, V. (2020). Machine learning applications for precision agriculture: A comprehensive review. *IEEE Access*, *9*, 4843–4873. https://doi.org/10.1109/ACCESS.2020.3048415

Sood, A., & Craven, M. (2022). Feature importance explanations for temporal black-box models. *Proceedings of the AAAI Conference on Artificial Intelligence*, *36*(8), 8351–8360. https://doi.org/10.1609/aaai.v36i8.20810

Teshome, F. T., Bayabil, H. K., Schaffer, B., Ampatzidis, Y., & Hoogenboom, G. (2024). Improving soil moisture prediction with deep learning and machine learning models. *Computers and Electronics in Agriculture*, *226*, 109414. https://doi.org/10.1016/j.compag.2024.109414

Thapa, V. R., Ghimire, R., VanLeeuwen, D., Acosta-Martínez, V., & Shukla, M. (2022). Response of soil organic matter to cover cropping in water-limited environments. *Geoderma*, *406*, 115497. https://doi.org/10.1016/j.geoderma.2021.115497

Wang, D., Thunéll, S., Lindberg, U., Jiang, L., Trygg, J., & Tysklind, M. (2022). Towards better process management in wastewater treatment plants: Process analytics based on SHAP values for tree-based machine learning methods. *Journal of Environmental Management*, *301*, 113941. https://doi.org/10.1016/j.jenvman.2021.113941

Wang, K. Z., Gao, M., Mukundan, R., Gelda, R., & Frei, A. (2025). Application of tree-based machine learning models to predict phosphorus levels in a water supply stream. *Journal of Hydroinformatics*, *27*(8), 1292–1308. https://doi.org/10.2166/hydro.2025.025

Wang, L. J., Shi, L. S., Reimers, C., Wang, Y. L., He, L. L., Wang, Y. J., Reichstein, M., & Jiang, S. J. (2025). A self-supervised deep learning model for enhanced generalization in soil moisture prediction. *Journal of Hydrology*, *662*, 133974. https://doi.org/10.1016/j.jhydrol.2025.133974

Wang, Y., Shi, L., Hu, Y., Hu, X., Song, W., & Wang, L. (2024). A comprehensive study of deep learning for soil moisture prediction. *Hydrology and Earth System Sciences*, *28*(4), 917–943. https://doi.org/10.5194/hess-28-917-2024

Xi, X., Zhuang, Q., & Liu, X. (2025). A hybrid physics-guided deep learning modeling framework for predicting surface soil moisture. *Journal of Geophysical Research: Machine Learning and Computation*, *2*(3), e2025JH000682. https://doi.org/10.1029/2025JH000682

Yang, S., Lian, H., Xu, B., Thanh, H. V., Chen, W., Yin, H., & Dai, Z. (2023). Application of robust deep learning models to predict mine water inflow: Implication for groundwater environment management. *Science of the Total Environment*, *871*, 162056. https://doi.org/10.1016/j.scitotenv.2023.162056

Yang, Y., Hu, J., Porter, D., Marek, T., Heflin, K., & Kong, H. (2020). Deep reinforcement learning-based irrigation scheduling. *Transactions of the ASABE*, *63*(3), 549–556. https://doi.org/10.13031/trans.13633

Yin, H., Carroll, K. C., Yuan, Y., Jamil, A., Rucker, D. F., Dai, Z., & Soltanian, M. R. (2025). Hybrid vision transformer with convolutional blocks approach for subsurface electrical resistivity tomography inversion. *Journal of Geophysical Research: Machine Learning and Computation*, *2*, e2025JH000734. https://doi.org/10.1029/2025JH000734

Yin, H., Ikard, S. J., Rucker, D. F., Brooks, S. C., Dai, Z., & Carroll, K. C. (2025). Imaging hyporheic exchange by integrating deep learning and physics-informed inversion of time-lapse self-potential data. *Geophysical Research Letters*, *52*, e2025GL118772. https://doi.org/10.1029/2025GL118772

Yin, H., Wu, Q., Yin, S., Dong, S., Dai, Z., & Soltanian, M. R. (2023). Predicting mine water inrush accidents based on water level anomalies of borehole groups using long short-term memory and isolation forest. *Journal of Hydrology*, *616*, 128813. https://doi.org/10.1016/j.jhydrol.2022.128813

Yin, H., Zhang, G., Wu, Q., Cui, F., Yan, B., Yin, S., Soltanian, M. R., Thanh, H. V., & Dai, Z. (2024). Transfer learning with transformer-based models for mine water inrush prediction: A multivariate analysis using sparse and imbalanced monitoring data. *Mine Water and the Environment*, *43*(4), 707–726. https://doi.org/10.1007/s10230-024-01011-2

Zhan, C., Dai, Z., Jiao, J. J., Soltanian, M. R., Yin, H., & Carroll, K. C. (2025). Toward artificial general intelligence in hydrogeological modeling with an integrated latent diffusion framework. *Geophysical Research Letters*, *52*, e2024GL114298. https://doi.org/10.1029/2024GL114298

Zhan, C., Dai, Z., Yin, S., Carroll, K. C., & Soltanian, M. R. (2024). Conceptualizing future groundwater models through a ternary framework of multisource data, human expertise, and machine intelligence. *Water Research*, *257*, 121679. https://doi.org/10.1016/j.watres.2024.121679

Zhang, L., Bai, G., Evett, S. R., Colaizzi, P. D., Xue, Q., Marek, G., Dhungel, R., Zhao, H., Wan, N., & Lin, X. (2025). Increased irrigation could mitigate future warming-induced maize yield losses in the Ogallala Aquifer. *Communications Earth & Environment*, *6*(1), Article 483. https://doi.org/10.1038/s43247-025-02459-y

Zhang, R. P., Yu, H. Y., Zhang, W. B., Li, W., Su, H., Wu, S. X., Xu, Q., Li, Y. Y., & Yao, H. Y. (2024). Straw return enhances grain yield and quality of three main crops: Evidence from a meta-analysis. *Frontiers in Plant Science*, *15*, 1433220. https://doi.org/10.3389/fpls.2024.1433220

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Yin, H., Bista, P., Ghimire, R., Yang, H., & Carroll, K. C. (2026). Unraveling soil moisture dynamics with dual-scale interpretable machine learning: Cover cropping and irrigation insights in semi-arid agriculture. *Vadose Zone Journal*, *25*, e70077. https://doi.org/10.1002/vzj2.70077